NVIDIA DLI COURSE CATALOG

July 2022

NVIDIA | DEEP LEARNING INSTITUTE

Introduction

The NVIDIA Deep Learning Institute (DLI) offers resources for diverse needs, giving individuals and teams, organizations, educators, and students what they need to advance their knowledge in AI, accelerated computing, data science, graphics and simulation, networking, and more.

With access to GPU-accelerated servers in the cloud, you will learn how to train, optimize, and deploy neural networks using the latest deep learning tools, frameworks, and SDKs. You will also learn how to assess, parallelize, optimize, and deploy GPU-accelerated computing applications.

The DLI offers both self-paced online courses and instructor-led, pre-scheduled workshops. The self-paced courses range from 10 to 60 minutes and guide you through setting up a project or applying a specific technology, exposing you to fundamental skills. The 8-hour instructor-led workshops go deeper into topic areas, teaching you how to implement a project end-to-end. Both types of courses give you valuable hands-on experience using the latest technologies.

Why Choose the NVIDIA Deep Learning Institute for Training?

- Learn to build deep learning and accelerated computing applications for industries such as healthcare, robotics, autonomous driving, manufacturing, and more.
- Gain hands-on experience with the most widely used, industry-standard software, tools and frameworks. Each student will have access to a fully configured, GPU-accelerated server in the cloud.
- Access instructor-led workshops and online courses from anywhere using just a laptop and
internet connection.
- Acquire real-world expertise through content designed in collaboration with industry leaders such as Children's Hospital of Los Angeles, Mayo Clinic, and PwC.
- Earn NVIDIA DLI course completion certificates to indicate subject matter competency and support your career growth.

CERTIFICATE
Participants can earn a certificate to demonstrate subject matter competency and support professional career growth. Certificates are offered for select instructor-led workshops and online courses.

INSTRUCTOR-LED WORKSHOPS

WORKSHOP TITLE / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	CERTIFICATE
Deep Learning				
Fundamentals of Deep Learning Learn how deep learning (DL) works through hands-on exercises in computer vision and natural language processing (NLP). You will train deep learning models from scratch, and pick up tricks and tools for achieving highly accurate results along the way. 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. >**Learn more**	An understanding of fundamental programming concepts in **Python 3**, such as functions, loops, dictionaries, and arrays. Also, familiarity with **pandas data structures** and an understanding of how to compute a **regression line**. Suggested materials to satisfy prerequisites: **Python Beginner's Guide**	TensorFlow 2 with Keras, pandas	English, Japanese, Korean, Simplified Chinese, Traditional Chinese	✓
Building Transformer-Based Natural Language Processing Explore how to use Transformer-based natural language processing models for text classification tasks, such as categorizing documents. You will also explore 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. >**Learn more**	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.	PyTorch, pandas, NVIDIA NeMo™, NVIDIA Triton™ Inference Server	English, Korean, Simplified Chinese	✓

WORKSHOP TITLE / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	CERTIFICATE
Building Conversational AI Applications Discover how to quickly build and deploy production quality conversational AI applications with real-time transcription and natural language processing capabilities. **>Learn more**	Experience with Python coding and use of library functions and parameters. Also, a fundamental understanding of a deep learning framework, such as TensorFlow, PyTorch, or Keras and a basic understanding of neural networks.	NVIDIA Riva, NVIDIA TAO Toolkit, Kubernetes	English	✓
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. **>Learn more**	Intermediate knowledge of Python, including an understanding of list comprehension. Data science experience using Python and familiarity with NumPy and matrix mathematics.	CuDF, CuPy, TensorFlow 2, NVIDIA Merlin™, NVTabular, and NVIDIA Triton Inference Server	English	✓
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. **>Learn more**	Experience with stochastic-gradient-descent mechanics, network architecture, and parallel computing	TensorFlow, Keras, Horovod	English, Korean, Japanese, Simplified Chinese, Traditional Chinese	✓

WORKSHOP TITLE / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	CERTIFICATE
Model Parallelism: Building and Deploying Large Neural Networks Very large deep neural networks (DNNs), whether applied to natural language processing (e.g., GPT-3), computer vision (e.g., huge Vision Transformers), or speech AI (e.g., Wave2Vec 2) have certain properties that set them apart from their smaller counterparts. As DNNs become larger and are trained on progressively larger data sets, they can adapt to new tasks with just a handful of training examples, accelerating the route toward general AI. Training models that contain tens to hundreds of billions of parameters on vast data sets isn't trivial and requires a unique combination of AI, high performance computing (HPC), and systems knowledge. The goal of this course is to demonstrate how to train the largest neural networks and deploy them to production. **>Learn more**	> Good understanding of **PyTorch**, **deep learning**, and **data parallel** training concepts. > Practice with **multi-GPU training** and **natural language processing** are useful, but optional	PyTorch, Megatron-LM, DeepSpeed, Slurm, NVIDIA Triton™ Inference Server, NVIDIA Nsight™	English	✓
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). **>Learn more**	Experience with CNNs and Python	NVIDIA RAPIDS™, XGBoost, TensorFlow, Keras, pandas, autoencoders, GANs	English	✓
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. **>Learn more**	Experience with Python and deep networks	Python, TensorFlow, Keras, XGBoost, RAPIDS, cuDF, long short-term memory (LSTM), autoencoders	English	✓

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	CERTIFICATE
Computer Vision for Industrial Inspection Deep learning-based computer vision technology enables manufacturers to perform automated visual inspection. Compared to traditional visual inspection processes—which are often manual and rules-based—visual inspection AI can improve efficiency, reduce operating costs, and deliver more consistent results. In this DLI workshop, developers will learn how to create an end-to-end hardware-accelerated industrial inspection pipeline to automate defect detection. Using NVIDIA's own production data set as an example, we will illustrate how the application can be easily applied to a variety of manufacturing use cases. Developers will also learn to identify and mitigate common pitfalls in deep learning-based computer vision tasks, and be able to deploy and measure the effectiveness of their AI solution. All workshop attendees get access to fully configured, GPU-accelerated servers in the cloud, guidance from a DLI certified instructor, and the opportunity to network with other developers, data scientists, and researchers attending the workshop. Attendees can also earn a certificate to prove subject matter competency and support professional growth. >**Learn more**	> Experience with Python; basic understanding of data processing and deep learning > To gain experience with Python, we suggest this **Python tutorial** > For a basic understanding of data processing and deep learning, we suggest the DLI's **Fundamentals of Deep Learning**	Python, pandas, DALI, NVIDIA TAO Toolkit, NVIDIA TensorRT™, and NVIDIA Triton Inference Server	English	✓

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	CERTIFICATE
Building AI-Based Cybersecurity Pipelines Traditional cybersecurity methods include creating barriers around your infrastructure to protect it from intruders. However, as enterprises continue to digitally transform, they're faced with a proliferation of devices, more sophisticated cybersecurity attacks, and an incredibly vast network of data to protect—which means new cybersecurity methodologies must be explored. An alternative approach is to address cybersecurity as a data science problem: Aim to better understand all the users and activities across your network so that you can identify which transactions are typical and which are potentially nefarious. The **NVIDIA Morpheus AI framework** lets cybersecurity developers and practitioners harness the power of GPU computing to implement cybersecurity solutions that perform on a scale never before possible. With Morpheus, cybersecurity developers can create optimized AI pipelines for filtering, processing, and classifying large volumes of real-time data. Bringing a new level of information security to data centers, Morpheus enables dynamic protection, real-time telemetry, and adaptive defenses for detecting and remediating cybersecurity threats. >**Learn more**	> Familiarity with defensive cybersecurity themes > Professional data science and/or data analysis experience > Competency with the Python programming language > Competency with the Linux command line	NVIDIA Morpheus, NVIDIA Triton™ Inference Server, NVIDIA RAPIDS™, CLX, Helm, Kubernetes	English	✓

WORKSHOP TITLE / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	CERTIFICATE
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 NVIDIA Nsight™ Systems profiler. >**Learn more**	Basic C/C++ competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. No previous knowledge of CUDA programming is assumed.	NVIDIA Nsight Systems, nsys	English, Korean, Japanese, Simplified Chinese, Traditional Chinese	✓
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. >**Learn more**	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.	CUDA, Python, Numba, NumPy	English, Simplified Chinese, Traditional Chinese	✓
Accelerating CUDA C++ Applications with Multiple GPUs Discover 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. >**Learn more**	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.	CUDA C++, NVCC, Nsight Systems	English, Simplified Chinese	✓

WORKSHOP TITLE / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	CERTIFICATE
Scaling CUDA C++ Applications to Multiple Nodes Learn the tools and techniques needed to write CUDA C++ applications that can scale efficiently to clusters of NVIDIA GPUs. >**Learn more**	Intermediate experience writing CUDA C/C++ applications	C++, CUDA, MPI, NVSHMEM	English	✓
Fundamentals of Accelerated Computing with OpenACC Find out how to write and configure code parallelization with OpenACC, optimize memory movements between the CPU and GPU accelerator, and apply the techniques to accelerate a CPU-only Laplace Heat Equation to achieve performance gains. >**Learn more**	Basic C/C++ or Fortran competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. No previous knowledge of GPU programming is assumed.	NVIDIA Nsight, OpenACC	English	✓

Data Science

Accelerating Data Engineering Pipelines Explore how to employ advanced data engineering tools and techniques with GPUs to significantly improve data engineering pipelines. >**Learn more**	Intermediate knowledge of **Python** (list comprehension, objects). Familiarity with **pandas** and **introductory statistics** (mean, median, mode) a plus.	pandas, cuDF, Dask, NVTabular, Plotly	English	✓

WORKSHOP TITLE / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	CERTIFICATE
Fundamentals of Accelerated Data Science 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. >**Learn more**	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.	RAPIDS, cuDF, XGBoost, cuML, cuGraph, Dask, cuPy, pandas, NumPy, Bokeh	English, Traditional Chinese	✓

Networking

The NVIDIA Academy offers customizable training and certification on dozens of networking topics, including InfiniBand, NVIDIA 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.nvidia.com**

SELF-PACED COURSES

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
Deep Learning						
Building a Brain in 10 Minutes This one-click notebook explores the biological and psychological inspirations for the world's first neural networks. >**Learn more**	An understanding of fundamental programming concepts in Python 3 such as functions, loops, dictionaries, and arrays.	N/A	English	10 minutes	Free	N/A
Getting Started with Deep Learning Explore the fundamentals of deep learning by training neural networks and using results to improve performance and capabilities. >**Learn more**	An understanding of fundamental programming concepts in **Python 3**, such as functions, loops, dictionaries, and arrays. Familiarity with **pandas data structures** and an understanding of how to compute a **regression line**. Suggested materials to satisfy prerequisites: **Python Beginner's Guide**	TensorFlow 2 with Keras, pandas	English	8 hours	$90 (excludes tax, if applicable)	✓
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. >**Learn more**	Basic familiarity with Python (helpful, not required)	PyTorch, Jetson Nano	English, Simplified Chinese	8 hours	Free (hardware required)	✓

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
Building Video AI Applications at the Edge on Jetson Nano Use JupyterLab notebooks and Python application samples on your Jetson Nano to build new projects that extract meaningful insights from video streams through deep learning video analytics. >**Learn more**	Basic familiarity with the Linux command line and an understanding of fundamental programming concepts in Python 3, such as functions, loops, dictionaries, and arrays	DeepStream, TensorRT, Jetson Nano, Python	English	8 hours	Free (hardware required)	N/A
Deploying a Model for Inference at Production Scale Learn how to deploy your own machine learning models on a GPU server. >**Learn more**	Familiarity with at least one machine learning framework, such as PyTorch, TensorFlow, ONNX, or TensorRT	NVIDIA Triton	English	4 hours	$30 (excludes tax, if applicable)	N/A
Getting Started with Image Segmentation Learn how to categorize segments of an image. >**Learn more**	Basic experience training neural networks	TensorFlow 2 with Keras	English	2 hours	$30 (excludes tax, if applicable)	N/A
Modeling Time-Series Data with Recurrent Neural Networks in Keras Explore how to classify and forecast time-series data using recurrent neural networks (RNNs), such as modeling a patient's health over time. >**Learn more**	Basic experience with deep learning	Keras	English	2 hours	$30 (excludes tax, if applicable)	N/A

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
Building Real-Time Video AI Applications Gain the knowledge and skills needed to enable the real-time transformation of raw video data from widely-deployed camera sensors into deep learning-based insights. **>Learn more**	Competency in the Python 3 programming language, some experience manipulating data using pandas DataFrames, and familiarity with deep networks (specifically variations of CNNs)	NVIDIA DeepStream, NVIDIA TAO Toolkit, and NVIDIA TensorRT	English	8 hours	$90.00 (excludes tax, if applicable)	N/A
Integrating Sensors with NVIDIA DRIVE® Find out how to integrate automotive sensors into your applications using NVIDIA DRIVE. **>Learn more**	Basic experience in C++ and Linux terminal commands	C++, NVIDIA DriveWorks	English	2 hours	$30 (excludes tax, if applicable)	N/A
Optimized Vehicle Routing NVIDIA cuOpt is a GPU-accelerated logistics solver that uses heuristics and optimizations to calculate complex vehicle-routing problems with a wide range of constraints. In this self-paced course, participants will work through a demonstration of a common vehicle-routing optimization problem. Upon completion, participants will be able to preprocess input data for use by NVIDIA cuOpt and compose variants of the problem that reflect real-world business constraints. **>Learn more**	Anyone can run the code to see how it works, but to get the most out of this content, we recommend: > An understanding of fundamental programming concepts in **Python 3** such as functions, loops, dictionaries, and arrays > A familiarity of matrix-based Python libraries such as **NumPy** and **pandas** > A familiarity with **NVIDIA RAPIDS**; in particular **CuDF**, is nice to have but not required	NVIDIA cuOpt, CuDF, SciPy, NumPy, pandas, GeoPandas, VeRoViz	English	1 hours	Free	N/A

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
Train a Text Classification Model Using the NVIDIA TAO Toolkit This course includes a sample notebook from the **TAO Toolkit - Conversational AI** collection in the NVIDIA NGC™ catalog, complete with a live GPU environment. The tutorial takes about 60 minutes to complete. In this tutorial, you'll learn how to leverage the simplicity and convenience of TAO (Train, Adapt, Optimize) to build a text classification task for sentiment analysis. You'll learn how to: > Train and fine-tune a BERT text classification model on the SST-2 dataset > Run evaluation and inference on the model > Export the model to an ONNX format or NVIDIA® Riva format for deployment >**Learn more**	This course requires that the user have an NGC account and API key. To fulfill this requirement, please: 1. **Register and activate a free NGC account** 2. **Generate your NGC API key** and save it in a safe location	**NVIDIA TAO Toolkit**, **NVIDIA NGC**	English	1 hours	Free	N/A

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
Deploy a Text Classification Model Using NVIDIA Riva This course includes a sample notebook from the **NVIDIA TAO Toolkit - Conversational AI** collection in the NVIDIA NGC™ catalog, complete with a live GPU environment. The tutorial takes about 30 minutes to complete. In this lab, you'll learn how to deploy a text classification model with the NVIDIA® Riva ServiceMaker framework. The process to aggregate all the necessary artifacts for Riva deployment to a target environment is quick and straightforward. Once the model is deployed in Riva, you can issue inference requests to the server. > Use Riva ServiceMaker to take a TAO-exported .riva model and convert it to an aggregated .rmir file for final deployment > Deploy the model(s) locally on the Riva server > Send inference requests from a demo client using Riva API bindings Upon completion, you'll be familiar with how to deploy a text classification model using Riva on NVIDIA GPUs. >**Learn more**	This course requires that the user have an NGC account and API key. To fulfill this requirement, please: 1. **Register and activate a free NGC account** 2. **Generate your NGC API key** and save it in a safe location This tutorial doesn't require you to write novel code, but to best understand some of its details, you should have familiarity with fundamental programming concepts in Python 3 such as classes, functions, and loops.	**NVIDIA Riva**, **NVIDIA NGC**	English	30 minutes	Free	N/A

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
NVIDIA Riva Speech API Demo This course includes two sample notebooks from the **NVIDIA® Riva Speech Clients** container, complete with a live GPU environment. The tutorial takes about 30 minutes to complete. In this lab, you'll learn how to interact with the Riva server to process various conversational AI requests. You'll learn how to: > Send audio to an automatic speech recognition (ASR) model and receive back text > Use natural language processing (NLP) models to transform text, classify text, and classify tokens > Send text to a text-to-speech (TTS) model and receive back audio Upon completion, you'll be familiar with how to construct requests to the Riva server from a sample client. >**Learn more**	This tutorial doesn't require you to write novel code, but to best understand some of its details, you should have familiarity with fundamental programming concepts in Python 3 such as classes, functions, and loops.	**NVIDIA Riva**	English	30 minutes	Free	N/A

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
Sensitive Information Detection with NVIDIA Morpheus The NVIDIA Morpheus AI framework lets cybersecurity developers and practitioners harness the power of GPU computing to implement cybersecurity solutions that perform on a scale never before possible. In this free tutorial, you'll do multi-class sensitive information detection. You'll use Morpheus to ingest and preprocess data, perform AI-enabled inference on the data, and stream the results in real time for analysis and action. >**Learn more**	This tutorial doesn't have any prerequisites, but familiarity with defensive cybersecurity themes and the Linux command line are a plus.	**NVIDIA Morpheus AI framework, NVIDIA Triton™ Inference Server**	English	1 hour	Free	N/A
Disaster Risk Monitoring Using Satellite Imagery Learn how to build and deploy a deep learning model to automate the detection of flood events using satellite imagery. This workflow can be applied to lower the cost, improve the efficiency, and significantly enhance the effectiveness of various natural disaster management use cases. >**Learn more**	> Competency in the **Python 3** programming language > Basic understanding of machine learning and deep learning concepts, specifically variations of convolutional neural networks (**CNNs**), and pipelines > Interest in understanding how to manipulate satellite imagery using modern methods.	NVIDIA DALI™, the NVIDIA TAO Toolkit, NVIDIA® TensorRT™, NVIDIA Triton™ Inference Server	English	4 hours	Free	N/A

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
Accelerated Computing Fundamentals						
An Even Easier Introduction to CUDA Learn the basics of writing parallel CUDA kernels to run on NVIDIA GPUs. >Learn more	Competency writing applications in CUDA C/C++	C/C++	English	1 hour	Free	N/A
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. >Learn more	Basic C/C++ competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. No previous knowledge of CUDA programming is assumed.	C/C++, CUDA	English, Japanese, Korean, Simplified Chinese, Traditional Chinese	8 hours	$90 (excludes tax, if applicable	✓
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. >Learn more	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.	CUDA, Python, Numba, NumPy	English, Simplified Chinese, Traditional Chinese	8 hours	$90 (excludes tax, if applicable	✓

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
Scaling Workloads Across Multiple GPUs with CUDA C++ Learn how to build robust and efficient CUDA C++ applications that can leverage all available GPUs on a single node. >**Learn more**	Competency writing applications in CUDA C/C++	C/C++, accelerated computing, CUDA	English	4 hours	$30 (excludes tax, if applicable)	✓
Accelerating CUDA C++ Applications with Concurrent Streams Discover how to improve performance for your CUDA C/C++ applications by overlapping memory transfers to and from the GPU with computations on the GPU. >**Learn more**	Competency writing applications in CUDA C/C++	C/C++	English	4 hours	$30 (excludes tax, if applicable)	✓

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
Optimizing CUDA Machine Learning Codes with Nsight Profiling Tools NVIDIA Developer Tools are a collection of applications, spanning desktop and mobile targets, which enable developers to build, debug, profile, and develop class-leading and cutting-edge software utilizing the latest visual computing hardware from NVIDIA. In this course, you will learn the effective use of two powerful NVIDIA developer tools: **Nsight Systems** and **Nsight Compute**. Nsight Systems provide developers with a system-wide visualization of an application's performance. Developers can optimize bottlenecks to scale efficiently across any number or size of CPU and GPU—from large servers to the smallest systems on chip. Nsight Compute is an interactive kernel profiler for CUDA applications. It provides detailed performance metrics and API debugging via a user interface and command line tool. By the time you complete this course, you will be able to use Nsight Systems and Nsight Compute to analyze and optimize CUDA applications. Following best practices, you will begin by using Nsight Systems to analyze overall application structure and explore parallelization opportunities before turning to Nsight Compute to analyze and optimize individual CUDA kernels. >**Learn more**	Familiarity with machine learning applications using CUDA. We suggest DLI's **Fundamentals of Accelerated Computing with CUDA C/C++**	**NVIDIA Nsight Systems**, **NVIDIA Nsight Compute**	English	2 hours	$30	N/A

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
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. >Learn more	Basic experience with C/C++	OpenACC, C/C++	English	8 hours	$90 (excludes tax, if applicable	N/A
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. >Learn more	Proficiency programming in C/C++ and professional experience working on HPC applications	Docker, Singularity, HPC Container Maker (HPCCM), C/C++	English	2 hours	$30 (excludes tax, if applicable)	N/A

Data Science

Speed Up DataFrame Operations with RAPIDS cuDF This one-click notebook demonstrates significant speed-up by moving common DataFrame operations to the GPU with minimal changes to existing code. >Learn more	An understanding of fundamental programming concepts in Python 3 such as functions, loops, dictionaries, and arrays.	N/A	English	20 minutes	Free	N/A

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
Accelerating End-to-End Data Science Workflows Explore 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. >**Learn more**	Experience with Python, ideally including pandas and NumPy	RAPIDS, cuDF, cuML, cuGraph, Apache Arrow	English	6 hours	$90 (excludes tax, if applicable	✓

Graphics and Simulation

Getting Started with USD for Collaborative 3D Workflows Learn how to generate a scene using human-readable Universal Scene Description ASCII (.USDA) files. Upon completion, participants will be able to create their own scenes within the USD framework and will have a strong foundation to use it in applications, such as **NVIDIA Omniverse™**, Maya, Unity, and Unreal Engine. >**Learn more**	A basic understanding of computer graphics concepts, such as vertices, meshes, and RGB values and an understanding of fundamental programming concepts in Python like functions, loops, dictionaries, and arrays.	**Universal Scene Description Framework**	English	2 hours	Free	N/A

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE

Infrastructure

Introduction to AI in the Data Center

Explore AI, GPU computing, NVIDIA AI software architectures, and how to implement and scale AI workloads in the enterprise data center. >**Learn more**	Basic knowledge of enterprise networking, storage, and data center operations	Artificial intelligence, machine learning, deep learning, GPU hardware and software	English	4 hours	$30 (excludes tax, if applicable)	Digital Badge Available

Introduction to DOCA for DPUs

The NVIDIA DOCA™ Software Framework enables developers to rapidly create applications and services on top of NVIDIA BlueField® data processing units (DPUs). Together, DOCA and the BlueField DPU deliver breakthrough networking, security, and storage performance with a comprehensive, open development platform. In this self-paced course, participants will learn the basic concepts of DOCA as a platform for accelerated data center computing on BlueField DPUs. Upon completion, participants will be equipped with introductory knowledge that will enable them to begin using DOCA and DPUs to develop applications that accelerate their data centers services. >**Learn more**	> Familiarity with software architecture and how it relates to and executes on hardware. Suggested materials to satisfy prerequisite: • **Enterprise Data Center Networking** • **Data Center - Overview** • **Data Center - Virtualization** > Some working knowledge of data center networking. Suggested materials to satisfy prerequisite: • **Introducing How Computers Works** • **Hardware Acceleration** • **Software Execution and Computing**	**NVIDIA DOCA SDK**	English	2 hours	Free	N/A

WORKSHOP NAME / DESCRIPTION	PREREQUISITES	TOOLS, LIBRARIES, FRAMEWORKS	LANGUAGES	DURATION	PRICE	CERTIFICATE
Getting Started with DOCA Flow NVIDIA DOCA is the key to unlocking the potential of the NVIDIA BlueField data processing unit (DPU), enabling you to offload, accelerate, and isolate data center workloads. With DOCA, developers can program the data center infrastructure of tomorrow by creating software-defined, cloud-native, DPU-accelerated services with zero-trust protection to address the increasing performance and security demands of modern data centers. DOCA Flow is the most fundamental API for building generic execution pipes in hardware. The library provides an API for building a set of pipes, where each pipe consists of match criteria, monitoring, and a set of actions. Pipes can be chained so that after a pipe-defined action is executed, the packet may proceed to another pipe. In this course, you will be introduced to DOCA Flow programming by building an "ARP Storm Control" application which prevents network failures caused by broadcast storms. It does so through the creation of a DOCA Flow pipeline that can dampen malicious broadcast network activity without impacting well-behaved traffic. >**Learn more**	A working knowledge of networking basics	**DOCA Flow**	English	2 hours	Free	N/A

Networking

The NVIDIA 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 what's available, visit **academy.nvidia.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