

Fundamentals of Deep Learning for Multiple Data Types

This workshop uses a series of hands-on exercises to teach deep learning techniques for a range of problems involving multiple data types. After a quick introduction to deep learning, you'll advance to building deep learning applications for image segmentation, sentence generation, and image and video captioning, while learning relevant computer vision, neural network, and natural language processing concepts. At the end of the workshop, you'll be able to assess a broad spectrum of problems where deep learning can be applied.

Duration:	8 hours
Price:	\$10,000 for groups of up to 20 (price increase for larger groups). During the workshop, each participant will have dedicated access to a fully configured, GPU accelerated workstation in the cloud.
Assessment type:	Multiple-choice
Certificate:	Upon successful completion of the assessment, participants will receive an NVIDIA DLI certificate to recognize their subject matter competency and support professional career growth.
Prerequisites:	Familiarity with basic Python (functions and variables); prior experience training neural networks
Languages:	English, Japanese, Korean, Chinese
Tools, libraries, and frameworks:	TensorFlow, TensorBoard

Learning Objectives

At the conclusion of the workshop, you'll have an understanding of the fundamentals of deep learning for multiple data types and be able to:

- > Implement common deep learning workflows such as image segmentation and text generation
- > Compare and contrast data types, workflows, and frameworks
- > Combine deep learning-powered computer vision and natural language processing to start solving sophisticated real-world problems that require multiple input data types

Why Deep Learning Institute Hands-On Training?

- > Learn to build deep learning and accelerated computing applications for industries such as autonomous vehicles, finance, game development, healthcare, robotics, and more.
- > Obtain hands-on experience with the most widely used, industry-standard software, tools, and frameworks.
- > 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 DLI certificate to demonstrate your subject matter competency and support career growth.
- > Access content anywhere, anytime with a fully configured, GPU-accelerated workstation in the cloud.

Workshop Outline

TOPIC	DESCRIPTION
Introduction (15 mins)	<ul style="list-style-type: none"> > Meet the instructor. > Create an account at courses.nvidia.com/join
Image Segmentation with TensorFlow (120 mins)	<ul style="list-style-type: none"> > Compare image segmentation to other computer vision problems. > Experiment with TensorFlow tools. > Implement effective metrics for assessing model performance.
Break (60 mins)	
Word Generation with TensorFlow (120 mins)	<ul style="list-style-type: none"> > Learn about natural language processing (NLP) and recurrent neural networks (RNNs). > Create network inputs from text data. > Test with new data and iterate to improve performance.
Break (15 mins)	
Image and Video Captioning (120 mins)	<ul style="list-style-type: none"> > Combine computer vision and natural language processing to describe scenes. > Learn to harness the functionality of convolutional neural networks (CNNs) and RNNs.
Final Review (15 mins)	<ul style="list-style-type: none"> > Review key learnings and wrap up questions. > Complete the assessment to earn a certificate. > Take the workshop survey.