Fundamentals of Deep Learning for Natural Language Processing

This workshop teaches deep learning techniques for understanding textual input using natural language processing (NLP) through a series of hands-on exercises. You'll learn techniques to train a neural network for text classification, build a linguistic style model to extract features from a given text document, and create a neural machine translation model for converting text from one language to another.

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 to a fully configured, GPU-accelerated workstation in the cloud.
Assessment type:	Code-based, 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:	Basic experience with neural networks and Python programming; familiarity with linguistics
Languages:	English, Chinese
Tools, libraries, and frameworks:	TensorFlow, Keras

Learning Objectives

At the conclusion of the workshop, you'll have an understanding of:

- Classical approaches to convert text to a machine-understandable representation
- Implementation and properties of distributed representations (embeddings)
- Methods to train machine translators from one language to another

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)	> Meet the instructor. > Create an account at **courses.nvidia.com/join** > Explore the importance of data representation for computers to understand language, as well as NLP challenges and how to tackle them with deep learning.
Word Embeddings (120 mins)	> Learn about distributed data representations, such as word embeddings, using the Word2Vec algorithm. Once trained, word embeddings can be used for text classification.
Break (60 minutes)	
Text Classification (120 mins)	> Build a linguistic style model to extract features from a given set of texts using embeddings. > Use text classification to determine the authors of an unknown set of documents.
Break (15 mins)	
Text Translation (120 mins)	> Create a neural machine translation model to convert text from one language to another. > Learn the basic technique to translate human-readable text to machine-readable format. > Use attention mechanisms to improve results—especially for long strings.
Final Review (15 mins)	> Review key learnings and wrap up questions. > Complete the assessment to earn a certificate. > Take the workshop survey.