



Fundamentals of Deep Learning for Natural Language Processing

The [NVIDIA Deep Learning Institute](#) (DLI) and [The Data Science Institute at Bar-Ilan University](#) invite you to attend a hands-on deep learning workshop on 25-02-2019 from 09:00-17:00 at room 202 Building 604, exclusively for verifiable academic students, staff, and researchers.

NVIDIA DLI offers hands-on training for developers, data scientists, and researchers looking to solve challenging problems with deep learning and accelerated computing.

About This Workshop:

This workshop teaches deep learning techniques for understanding textual input using natural language processing (NLP) through a series of hands-on exercises. You will work with widely-used deep learning tools, frameworks, and workflows to perform neural network training on a fully-configured, GPU-accelerated workstation in the cloud. The course teaches 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.

Workshop Agenda:

Overview of Natural Language Processing (45 mins)	<ul style="list-style-type: none"> Importance of data representation for computers to understand language 	Overview of NLP challenges and how to tackle them with deep learning
Break (15 mins)		
Word Embeddings (120 mins)	<ul style="list-style-type: none"> Overview of word2vec algorithm for text classification 	We will cover distributed data representations, such as word embeddings using the word2vec algorithm. Once trained, the word embeddings can be used for variety of problems, including text classification.
Break (60 mins)		
Text Classification (120 mins)	<ul style="list-style-type: none"> Build a linguistic style model to extract features from a given set of texts using embeddings 	Text classification will be used to determine the authors of an unknown set of documents. The trained text-classification model is then used to identify the right author for a given text document.
Break (15 mins)		
Text Translation (120 mins)	<ul style="list-style-type: none"> 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, and how to use attention mechanisms to improve results — especially for long strings.
Closing Comments and Questions (15 mins)	<ul style="list-style-type: none"> Wrap-up, potential next steps, and Q&A 	Quick overview of the next steps you could leverage to build and deploy your own applications

Prerequisites:

Basic experience with neural networks and Python programming, familiarity with linguistics

Workshop Instructor:

Vered Shwartz

Workshop Setup Instructions:

1. Create an NVIDIA Developer account at <http://courses.nvidia.com/join>.
2. Make sure that WebSockets works for you:
 - Test your laptop at <http://websocketstest.com>
 - Under ENVIRONMENT, confirm that “WebSockets” is checked yes.
 - Under WEBSOCKETS (PORT 80), confirm that “Data Receive,” “Send,” and “Echo Test” are checked yes.
3. If there are issues with WebSockets, try updating your browser. We recommend Chrome, Firefox, or Safari for an optimal performance.
4. Once onsite, visit <http://courses.nvidia.com/dli-event> and enter the event code provided by the instructor.

[REGISTER NOW](#)

contact dsi@biu.ac.il with any questions regarding the tutorial event

This workshop is brought to you by:



DEEP
LEARNING
INSTITUTE



אוניברסיטת בר-אילן
Bar-Ilan University

