

## Building Transformer-Based Natural Language Processing Applications

Applications for natural language processing (NLP) have exploded in the past decade. With the proliferation of AI assistants and organizations infusing their businesses with more interactive human-machine experiences, understanding how NLP techniques can be used to manipulate, analyze, and generate text-based data is essential. Modern techniques can capture the nuance, context, and sophistication of language, just as humans do. And when designed correctly, developers can use these techniques to build powerful NLP applications that provide natural and seamless human-computer interactions within chatbots, AI voice agents, and more.

Deep learning models have gained widespread popularity for NLP because of their ability to accurately generalize over a range of contexts and languages. Transformer-based models, such as Bidirectional Encoder Representations from Transformers (BERT), have revolutionized NLP by offering accuracy comparable to human baselines on benchmarks like SQuAD for question-answer, entity recognition, intent recognition, sentiment analysis, and more.

In this workshop, you'll learn how to use Transformer-based natural language processing models for text classification tasks, such as categorizing documents. You'll also learn how to leverage Transformer-based models for named-entity recognition (NER) tasks and how to 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.

### Learning Objectives

By participating in this workshop, you'll be able to:

- > Understand how text embeddings have rapidly evolved in NLP tasks such as Word2Vec, recurrent neural network (RNN)-based embeddings, and Transformers
- > See how Transformer architecture features, especially self-attention, are used to create language models without RNNs
- > Use self-supervision to improve the Transformer architecture in BERT, Megatron, and other variants for superior NLP results
- > Leverage pre-trained, modern NLP models to solve multiple tasks such as text classification, NER, and question answering
- > Manage inference challenges and deploy refined models for live applications


## Workshop Information and Prerequisites:

<b>Duration:</b>	8 hours
<b>Price:</b>	<b>Contact us for pricing.</b>
<b>Prerequisites:</b>	<ul style="list-style-type: none"><li>&gt; Experience with Python coding and use of library functions and parameters</li><li>&gt; Fundamental understanding of a deep learning framework such as TensorFlow, PyTorch, or Keras</li><li>&gt; Basic understanding of neural networks</li></ul> <p>Suggested materials to satisfy prerequisites: <b>Python tutorial, Overview of Deep Learning Frameworks, PyTorch tutorial, Deep Learning in a Nutshell, Deep Learning Demystified.</b></p>
<b>Tools, libraries, and frameworks:</b>	PyTorch, pandas, NVIDIA NeMo™, NVIDIA Triton™ Inference Server
<b>Assessment type:</b>	<ul style="list-style-type: none"><li>&gt; Skills-based coding assessments evaluate students' ability to build an NLP task, including a neural module pipeline and training.</li><li>&gt; Multiple-choice questions evaluate students' understanding of the NLP concepts presented in the class.</li></ul>
<b>Certificate:</b>	Upon successful completion of the assessment, participants will receive an NVIDIA DLI certificate to recognize their subject matter competency and support professional career growth.
<b>Hardware Requirements:</b>	Desktop or laptop computer capable of running the latest version of Chrome or Firefox. Each participant will be provided with dedicated access to a fully configured, GPU-accelerated workstation in the cloud.
<b>Language:</b>	English

## Sample Workshop Outline

TOPIC	DESCRIPTION
<b>Introduction</b> (15 mins)	<ul style="list-style-type: none"> <li>&gt; Meet the instructor.</li> <li>&gt; Create an account at <a href="https://courses.nvidia.com/join">courses.nvidia.com/join</a></li> </ul>
<b>Introduction to Transformers</b> (120 mins)	<b>Explore how the Transformer architecture works in detail:</b> <ul style="list-style-type: none"> <li>&gt; Build the Transformer architecture in PyTorch.</li> <li>&gt; Calculate the self-attention matrix.</li> <li>&gt; Translate English to German with a pre-trained Transformer model.</li> </ul>
<b>Break</b> (60 mins)	
<b>Self-Supervision, BERT, and Beyond</b> (120 mins)	<b>Learn how to apply self-supervised Transformer-based models to concrete NLP tasks using NVIDIA NeMo:</b> <ul style="list-style-type: none"> <li>&gt; Build a text classification project to classify abstracts.</li> <li>&gt; Build a named-entity recognition (NER) project to identify disease names in text.</li> <li>&gt; Improve project accuracy with domain-specific models.</li> </ul>
<b>Break</b> (15 mins)	
<b>Inference and Deployment for NLP</b> (120 mins)	<b>Learn how to deploy an NLP project for live inference on NVIDIA Triton:</b> <ul style="list-style-type: none"> <li>&gt; Prepare the model for deployment.</li> <li>&gt; Optimize the model with NVIDIA® TensorRT™.</li> <li>&gt; Deploy the model and test it.</li> </ul>
<b>Final Review and Next Steps</b> (15 mins)	<ul style="list-style-type: none"> <li>&gt; Review key learnings and answer questions.</li> <li>&gt; Complete the assessment and earn a certificate.</li> <li>&gt; Take the workshop survey.</li> <li>&gt; Learn how to set up your own environment and discuss additional resources and training.</li> </ul>

### Why Choose NVIDIA Deep Learning Institute for Hands-On Training?

- > Access workshops from anywhere with just your desktop/laptop computer and an internet connection. Each participant will have access to a fully configured, GPU-accelerated workstation in the cloud.
- > Obtain hands-on experience with the most widely used, industry-standard software, tools, and frameworks.
- > Learn to build deep learning and accelerated computing applications for industries, such as healthcare, robotics, manufacturing, accelerated computing, and more.
- > 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 Deep Learning Institute certificate to demonstrate your subject matter competency and support your career growth. 

For the latest DLI Workshops and Trainings, visit [www.nvidia.com/dli](https://www.nvidia.com/dli).

For questions, contact us at [nvdl@nvidia.com](mailto:nvdl@nvidia.com).

© 2020 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, NeMo, TensorRT, and Triton are trademarks and/or registered trademarks of NVIDIA Corporation. All company and product names are trademarks or registered trademarks of the respective owners with which they are associated. Features, pricing, availability, and specifications are all subject to change without notice. AUG20