

## AI for Natural Language Processing (NLP)

### Course Summary

#### Description

We live in an era of so much data – a lot of it is text (emails, tweets, customer tickets, Yelp reviews, product reviews, etc.)

In the field of AI, there is a revolution going on in the past few years. Researchers from companies like Google, Facebook, Microsoft, and Baidu has come up with breakthrough algorithms that can understand text data more than ever before.

The applications are wide-ranging, including understanding documents, processing customer service tickets, and analyzing reviews.

In this course, we will teach how to handle text data and introduce you to modern AI NLP technologies.

#### Objectives

At the end of this course, students will be able to understand:

- How to prepare text for machine learning
- Stemming, tokenizing, and filtering stop words in text
- Analyzing documents using word-frequency, bag-of-words techniques
- Visualizing text data
- Classic toolsets for text processing: NLTK, Textblob, TF-IDF
- Naive Bayes for text classifications
- Modern techniques for text: Spacy, Word2Vec
- Topic modeling with Gensim
- Neural Networks and Deep Learning
- Deep learning models for text processing: LSTM, RNN
- Transformer architecture
- Modern NN models for text processing: ELMO, ULMFIT, BERT
- Text generation with Tensorflow

#### Topics

- Text Preparation
- AI overview
- AI with TensorFlow and Keras
- NLP and Deep Learning
- Unsupervised NLP (Optional)
- Conversational AI (Optional)

#### Audience

This course is designed for Developers, Data analysts, and Data Scientists.

#### Prerequisites

- Programming background
- Basic knowledge of Python language and Jupyter notebooks is recommended.  
Even if you haven't done any Python programming, Python is such an easy language to learn quickly. We will provide Python resources.

#### Duration

Three Days

## AI for Natural Language Processing (NLP)

### Course Outline

- I. *Machine Learning Overview*
  - A. Machine Learning landscape
  - B. Understanding AI use cases
  - C. Data and AI
  - D. AI vocabulary
  - E. Hardware and software ecosystem
  - F. Understanding types of Machine Learning (Supervised / Unsupervised / Reinforcement)
- II. *Text Preparation*
  - A. Filtering
  - B. Stopwords
  - C. Stemming
  - D. Parsing and tokenization
  - E. Word-clouds
  - F. Working with Unicode
- III. *Text Algorithms Overview*
  - A. N-grams
  - B. Bag-of-words
  - C. NLTK
  - D. TextBlob
  - E. TF-IDF
- IV. *Text Classification*
  - A. Naive Bayes
  - B. SVM
- V. *Text datasets and Benchmarks*
  - A. Public text datasets
  - B. Benchmarks (GLUE, SQUAD)
- VI. *Topic Modeling*
  - A. LDA (Latent Dirichlet Allocation)
  - B. Gensim
- VII. *Introduction to Neural Networks*
  - A. Perceptrons
  - B. Feedforward networks
  - C. Activation functions
  - D. Optimizers
  - E. Backpropagation
  - F. Deep Neural Networks
- VIII. *Tensorflow*
  - A. TensorFlow intro
  - B. TensorFlow features
  - C. TensorFlow on GPU and TPU
  - D. TensorFlow API
  - E. Lab: Setting up and Running TensorFlow
- IX. *NLP and Deep Learning*
  - A. Word embeddings
  - B. Skipgram
  - C. Training the model
  - D. Visualizing the embeddings
  - E. Word2Vec
  - F. SpaCy for named entity recognition
- X. *Recurrent Neural Networks (RNN)*
  - A. Introduction to RNNs
  - B. Text prediction
  - C. Named entity extraction
  - D. Automatic translation (seq2seq)
  - E. Text generation
- XI. *Transformers*
  - A. Attention concept
  - B. Transformer architecture
  - C. Bidirectional LSTM
  - D. Pre-trained Models for Text Processing (EIMO, ULMFIT, BERT)
- XII. *Conversational AI*
  - A. Understanding natural language
  - B. Generating natural language
  - C. Introduction to RASA framework
- XIII. *Final Workshop (Time Permitting)*
  - A. This a group exercise
  - B. Students will use the learned techniques to solve a real world problem
  - C. And present their solutions to the class
  - D. Discussions and Takeaways