Course Outline

Artificial Intelligence (AI) For Natural Language Processing (NLP) – Advanced

Course Summary

Description

We live in an era of so much data – and 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. The researchers from companies like Google, Facebook, Microsoft and Baidu have come up with break-through technologies that work with text based on its meaning and not the keywords in it.

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

In this course, we will give you practical skill in handling text with modern AI NLP technologies.

Topics

- Modern techniques for text: Spacy, Word2Vec
- Topic modeling with Gensim
- Neural Network frameworks: Tensorflow & Keras
- NN models for text processing: LSTM, RNN
- Modern NN models for text processing: ELMO, ULMFIT, BERT
- Practical skills in building BERT model
- Hugging Face Library

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

Course Outline

Artificial Intelligence (AI) For Natural Language Processing (NLP) –Advanced

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 datasets and Benchmarks
- A. Public text datasets
- B. Benchmarks (GLUE, SQUAD)

III. Topic Modeling
- A. LDA (Latent Dirichlet Allocation)
- B. Gensim

IV. Introduction to Neural Networks
- A. Perceptrons
- B. Feedforward networks
- C. Activation functions
- D. Optimizers
- E. Backpropagation
- F.Deep Neural Networks

V. Tensorflow
- A. TensorFlow intro
- B. TensorFlow features
- C. TensorFlow on GPU and TPU
- D. TensorFlow API
- E. Lab: Setting up and Running TensorFlow

VI. 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

VII. Recurrent Neural Networks (RNN)
- A. Introduction to RNNs
- B. Text prediction
- C. Named entity extraction
- D. Automatic translation (seq2seq)
- E. Text generation

VIII. Transformers
- A. Attention concept
- B. Transformer architecture
- C. Bidirectional LSTM
- D. Pre-trained Models for Text Processing (ElMO, ULMFIT, BERT)

IX. BERT
- A. BERT model
- B. BERT use cases
- C. Legal
- D. Financial
- E. Medical
- F.BERT implemented in Hugging Face

X. Conversational AI
- A. Understanding natural language
- B. Generating natural language
- C. Introduction to RASA framework

XI. 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