



NLP using deep learning

40 academic hours

Overview

Natural language processing (NLP) or computational linguistics is one of the most important technologies of the information age. Applications of NLP are everywhere because people communicate almost everything in language: web search, advertising, emails, customer service, language translation, virtual agents, medical reports, etc. In recent years, deep approaches have obtained very high performance across many different NLP tasks, using single end-to-end neural models that do not require traditional, task-specific feature engineering. In this course, students will gain a thorough introduction to cutting-edge research in Deep Learning for NLP.

Who should attend?

Junior data scientists who wish to deepen their knowledge in Deep Learning and NLP

Prerequisites

- Proficiency in Python
- Basic knowledge in Linear algebra and calculus
- Basic probability and statistics.
- Foundations of Machine Learning

After attending the course attendees should be able to:

- Use deep learning frameworks to define and solve various of NLP tasks
- Read an NLP paper to keep up with future advancements.

Syllabus

- Deep learning using pytorch
 - ML basics
 - The training loop + linear model
 - Feed forward NN
 - pytorch utils
- NLP
 - Tasks:
 - Text classification
 - Language models
 - Question answering
 - Captioning
 - Summarization
 - Algorithms and architectures:
 - Word2Vec and GloVe
 - Language modeling and RNNs.
 - Seq2Seq
 - ConvNets
 - (self) attention
 - Papers:
 - Attention Is All You Need
 - BERT – Bidirectional Transformers
 - GPT3

Module 1- Deep learning using pytorch

- ML basics with numpy
- The training loop + the linear model - Introduction to torch
- Feedforward neural network - torch.nn.Sequential
- Regularization and optimization - torch.nn.Module
- Convolutional neural networks - utils.Dataset and utils.DataLoader

Module 2 – NLP

- Introduction to NLP
- Pre-deep learning NLP - WordNet and NLTK, BOW
- Word2Vec and GloVe - Introduction to genism
- Language modeling and RNNs - RNNs in pytorch
- Fancier versions of RNN + other tricks - [Bi]LSTM, GRU in pytorch.
- Conditional Language model (Seq2Seq + Attention + Translation) - Seq2Seq and attention in pytorch
- ConvNets for text representation - Intro to huggingface
- Question Answering models - huggingface datasets
- Attention is all you need + BERT + self attention
- Few shots in Deep learning + GPT3

Covering all topics and tasks.

Some (not much) dependency parsing, sentence classification.

The idea of word embedding, use-cases.

Language modeling (Neural and statistical) and how to evaluate it, vanilla RNN, use-cases.

Multilayers RNNs, [bi]LSTM, GRU, intro to Seq2Seq, sentence encoding.

Seq2Seq (sentence decoding), Attention (general idea), beam search, BLEU.

Includes homework - develop a char-cnn.

The SOTA techniques for question answering.

Covering the paper “Attention is All You Need” and BERT.

Covering the topic of GPT as a language model.