

An Introduction to Deep Learning for Natural Language Processing

Wednesday 15 May 2019 13:30 (1h 30m)

We will provide an introduction to the application of Artificial Neural Networks to a variety of tasks in Natural Language Processing (NLP), i.e. enabling computers to 'make sense' of human language. The tutorial will cover example problems in document classification, sentiment analysis, and sequence labeling; we will introduce common architecture variants, such as deep feed-forward networks, convolutional neural networks, and recurrent neural networks. A basic recollection of linear algebra and Python syntax may be useful, but there should be something for everyone in this tutorial either way. We have yet to decide whether there will be time and opportunity for hands-on exercises.

Primary authors: KUTUZOV, Andrei (University of Oslo); OEPEN, Stephan (University of Oslo)

Presenters: KUTUZOV, Andrei (University of Oslo); OEPEN, Stephan (University of Oslo)

Session Classification: Workshops II