



The Hebrew University of Jerusalem

Syllabus

DEEP LEARNING FOR NATURAL LANGUAGE PROCESSING - 67583

Last update 30-01-2020

HU Credits: 1

Degree/Cycle: 2nd degree (Master)

Responsible Department: Computer Sciences

Academic year: 0

Semester: 2nd Semester

Teaching Languages: English

Campus: E. Safra

Course/Module Coordinator: Dr. Omri Abend

Coordinator Email: omri.abend@mail.huji.ac.il

Coordinator Office Hours:

Teaching Staff:

Ms. Anna Rumshisky

Course/Module description:

Deep neural network models have become the go-to choice for many natural language processing problems, improving the state-of-the-art on a variety of tasks from machine translation and question answering to inference and dialogue generation. This course will provide a basic introduction to deep learning methods for natural language processing. Covered topics will include vector space lexical embedding models, recurrent neural networks and their use for language modeling, encoder/decoder sequence-to-sequence and attention-based architectures. We will discuss how these methods are used for representation learning and language generation, and consider some practical applications such as question answering and conversational agents.

Course/Module aims:

Course aims to provide an introduction to the modern deep learning techniques for natural language processing.

Learning outcomes - On successful completion of this module, students should be able to:

Understand the computational models used to process natural language. Build, train and deploy neural network computational models for text processing tasks such as text generation or classification.

Attendance requirements(%):

100

Teaching arrangement and method of instruction: The course will include a combination of lectures, hands-on tutorials and programming assignments.

Programming assignments and tutorials will be in Python, using PyTorch deep learning library. We will use Jupyter notebooks for coding assignments.

Course/Module Content:

Review of neural networks models. Lexical embedding models: count-based vs. predicted word vectors. Building a word embedding model. Recurrent neural networks. Training with backpropagation. Common loss functions.

Dropout and other regularization methods. Gated cell memory architectures (LSTMs/GRUs). Neural language models. Conditional language models. Sequence-to-sequence encoder/decoder architectures. Building a sequence-to-sequence encoder/decoder model. Seq2seq models with attention. Neural attention models for machine translation. Attention-only encoder/decoder architectures. Transformers. Contextualized lexical embedding models. ELMo, BERT.

Required Reading:

There is no required textbook. Readings will be distributed by instructor.

Additional Reading Material:

J. Eisenstein. Natural Language Processing. MIT Press.

Grading Scheme:

Additional information: