EPFL

EE-608	**Deep Learning For Natural Language Processing**

Henderson James

Cursus	Sem.	Type
Electrical Engineering		Opt.

Language of teaching	English
Credits	4
Session	
Exam	Multiple
Workload	120h
Hours	**56**
Courses	28
TP	28
Number of positions	**40**

Frequency

Every 2 years

Remark

Next time: Fall 2023

Summary

The Deep Learning for NLP course provides an overview of neural network based methods applied to text. The focus is on models particularly suited to the properties of human language, such as categorical, unbounded, and structured representations, and very large input and output vocabularies.

Content

Models
- Word embeddings
- LSTMs and CNNs for text
- Attention models
- Sequence-to-sequence models
- NN integration with decoding
- Multi-task learning

Applications
- Language modelling
- Machine translation
- Syntactic parsing
- Semantic parsing
- Dialogue systems

Keywords

Machine Learning, Natural Language Processing, Neural Networks.

Learning Prerequisites

Required courses

Undergraduate level probability, linear algebra, and programming.

Recommended courses

Courses on Machine Learning, Natural Language Processing (Human Language Technology, Computational Linguistics), or Artificial Intelligence would be useful.

Learning Outcomes

By the end of the course, the student must be able to:

- Identify appropriate deep learning architectures for different natural language processing tasks.
- Apply appropriate training and evaluation methodology to such models on large datasets using existing packages.

Assessment methods

Multiple.