

EE-608

Deep Learning For Natural Language Processing

Henderson James

Cursus	Sem.	Type
Computational and Quantitative Biology		Opt.
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 2027

Summary

This course covers advanced topics in deep learning architectures for natural language processing. The focus is on attention-based architectures, structure processing and variational-Bayesian approaches, and why these models are particularly suited to the properties of human language.

Content**Models and Methods**

- * Attention and Transformers
- * Sequence-to-sequence and graph-to-graph models
- * Multi-task and transfer learning
- * Variational-Bayesian models and Diffusion
- * LLMs, prompting, and RAG

Applications

- * Language modelling
- * Machine translation
- * Text generation
- * Knowledge extraction and embedding
- * Dialogue systems

Keywords

Machine Learning, Natural Language Processing, Large Language Models.

Learning Prerequisites**Required courses**

- * Introductory course on Machine Learning.
- * Undergraduate level probability, linear algebra, and programming.

Recommended courses

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

Learning Outcomes

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

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

Assessment methods

Multiple.

Resources

Moodle Link

- <https://go.epfl.ch/EE-608>