

CS-552

Modern natural language processing

Bosselut Antoine

Cursus	Sem.	Type
Computational and Quantitative Biology		Opt.
Computer and Communication Sciences		Opt.
Computer science minor	E	Opt.
Computer science	MA2, MA4	Opt.
Data Science	MA2, MA4	Obl.
Data science minor	E	Opt.
Digital Humanities	MA2, MA4	Opt.
Life Sciences Engineering	MA2, MA4	Opt.
SC master EPFL	MA2, MA4	Opt.

Language of teaching	English
Credits	8
Session	Summer
Semester	Spring
Exam	During the semester
Workload	240h
Weeks	14
Hours	6 weekly
Courses	3 weekly
Exercises	1 weekly
Project	2 weekly
Number of positions	

Summary

Natural language processing is ubiquitous in modern intelligent technologies, serving as a foundation for language translators, virtual assistants, search engines, and many more. In this course, students will learn algorithmic tools for tackling problems in modern NLP.

Content

This course includes lectures, exercises, a midterm exam, and a project. In lectures, we will cover the foundations of modern methods for natural language processing, such as word embeddings, recurrent neural networks, transformers, pretraining, and how they can be applied to important tasks in the field, such as machine translation and text classification. We will also cover issues with these state-of-the-art approaches (such as robustness, interpretability, sensitivity), identify their failure modes in different NLP applications, and discuss analysis and mitigation techniques for these issues.

In assignments, students will be evaluated on their ability to implement methods learned in class on closed-form problems developed by the course staff. In their project, students will be expected to apply techniques learned in lecture to a problem of the course staff's choosing. They will formulate the problem as an NLP task, propose a suitable evaluation to measure their progress, develop a model to solve the task, and provide analysis of the strengths and weaknesses of their method.

This course is of interest to MS / PhD student interested in modern methods and issues in natural language processing, both from a research and applied perspective.

Learning Prerequisites**Recommended courses**

- CS-233a or CS-233b Introduction to machine learning
- CS-433 Machine learning

Important concepts to start the course

- Python programming
- Probability and Statistics
- Linear Algebra
- Machine Learning concepts

Learning Outcomes

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

- Define basic problems and tasks in natural language processing (e.g., machine translation, summarization, text classification, language generation, sequence labeling, information extraction, question answering)
- Implement common modern approaches for tackling NLP problems and tasks (embeddings, recurrent neural models, attentive neural models) and how to train them
- Understand failure modes of these models and learning algorithms (e.g., robustness, interpretability/explainability, bias, evaluation)
- Review academic research papers and understand their contributions, strengths, and weaknesses according to the principles learned in lecture
- Complete a project that applies these algorithms to a real-world NLP problem, where they will define a task, evaluation, model implementation, and analyze the shortcomings of their approach

Teaching methods

- Lectures
- Lab sessions
- Midterm Examination
- Course project

Expected student activities

- Attend lectures and participate in class
- Complete homework assignments
- Pass midterm exam
- Complete a project (set by the course supervisor) : complete a project proposal outlining topic and evaluation plan; submit two project milestones; submit final project report; present project findings to committee of course instructor and TAs.

Assessment methods

- Midterm examination (30%)
- Project (70%)

Supervision

Office hours	Yes
Assistants	Yes
Forum	Yes

Resources

Moodle Link

- <https://go.epfl.ch/CS-552>