

# CS-552: Modern NLP

## Introduction

Antoine Bosselut

EPFL



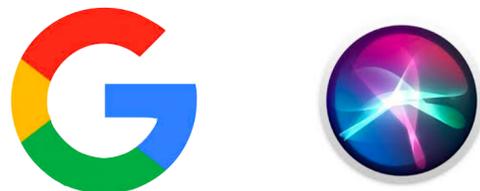
# Natural Language Processing

## Enabling Human-Machine Collaboration

Search Engines

Chatbots

Assistants



## Accelerating Human-Human Communication

Machine Translation

Text Summarization

Information Extraction



## Mining Human Insights

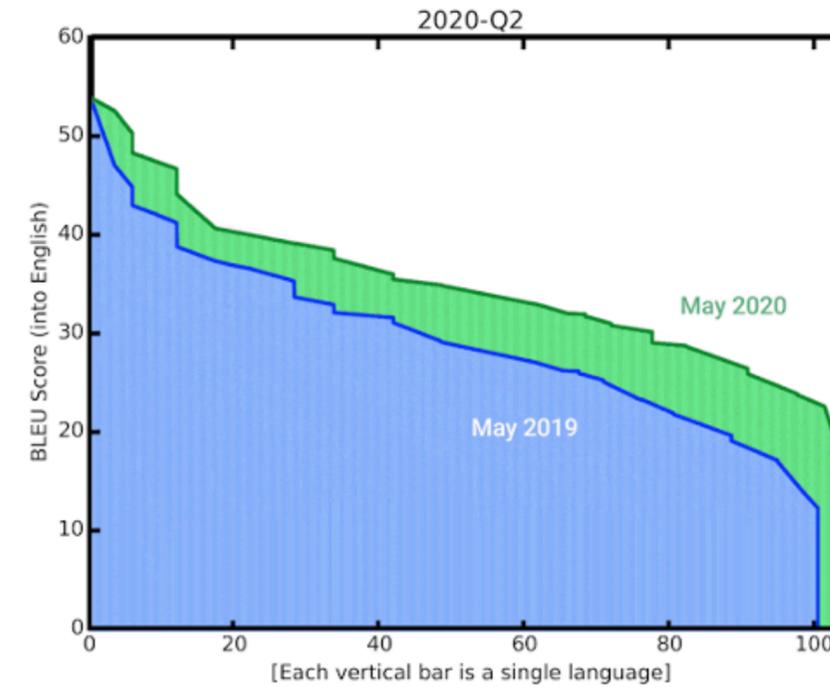
Sentiment Analysis

Motivation Analysis

Emotion Detection



# Machine Translation



DETECT LANGUAGE **FRENCH** ENGLISH ▼ ↔ **ENGLISH** FRENCH SPANISH ▼

J'ai mangé avec mon avocat ×  
aujourd'hui

I ate with my lawyer today ☆

38 / 5000

The screenshot shows the Google Translate interface. The source language is French and the target language is English. The source text is "J'ai mangé avec mon avocat aujourd'hui" and the translated text is "I ate with my lawyer today". The interface includes a microphone icon, a speaker icon, a character count of 38 / 5000, and a keyboard icon. The translated text has a star icon for saving and icons for copy, share, and refresh.

# Question Answering

what is the tallest mountain in europe ?



All Images Maps News Videos More Tools

About 12'400'000 results (1.10 seconds)

Europe / Mountains / Maximal / Elevation

## Mount Elbrus

5,642 m

### Mountains



Mount Elbrus  
5,642 m



Dykh-Tau  
5,205 m



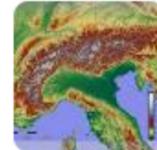
Shkhara  
5,193 m



Koshtan-Tau  
5,151 m



Mount Kazbek  
5,033 m



Alps



Mont Blanc  
4,810 m

who is the president of switzerland ?



All Images News Maps Videos More Tools

About 415'000'000 results (0.82 seconds)

## Guy Parmelin

The President of the Swiss Confederation in 2021 is **Guy Parmelin** from the canton of Vaud. He was elected on 9 December 2020. The President's department in 2021 is the Federal Department of Economic Affairs, Education and Research EAER.

<https://www.admin.ch> > gov > start > federal-presidency

[Presidential year 2021](#)

About featured snippets • Feedback

# Chatbots

AN

Can you give me a two-day itinerary for a weekend trip to Paris ? Please don't include more than three activities per day.



Day 1:

1. Visit the Eiffel Tower - Start your day with the iconic symbol of Paris, the Eiffel Tower. Take a stroll around the Champ de Mars and enjoy the views from the tower.
2. Explore the Louvre Museum - After lunch, head to the Louvre Museum and admire the world-famous artwork such as the Mona Lisa and the Venus de Milo.
3. Stroll along the Seine River - In the evening, take a romantic stroll along the Seine River and enjoy the Parisian atmosphere.

Day 2:

1. Visit Notre-Dame Cathedral - Start your day by visiting one of the most famous Gothic cathedrals in the world, Notre-Dame Cathedral.
2. Explore the Marais neighborhood - After lunch, explore the Marais neighborhood, known for its historical architecture, trendy boutiques, and restaurants.
3. Enjoy a Cabaret Show - In the evening, enjoy a cabaret show at the famous Parisian venue, the Moulin Rouge, and end your trip with a memorable experience.

**Why is natural language processing difficult?**

# Ambiguity

Kids make nutritious snacks

Stolen painting found by tree

Miners refuse to work after death

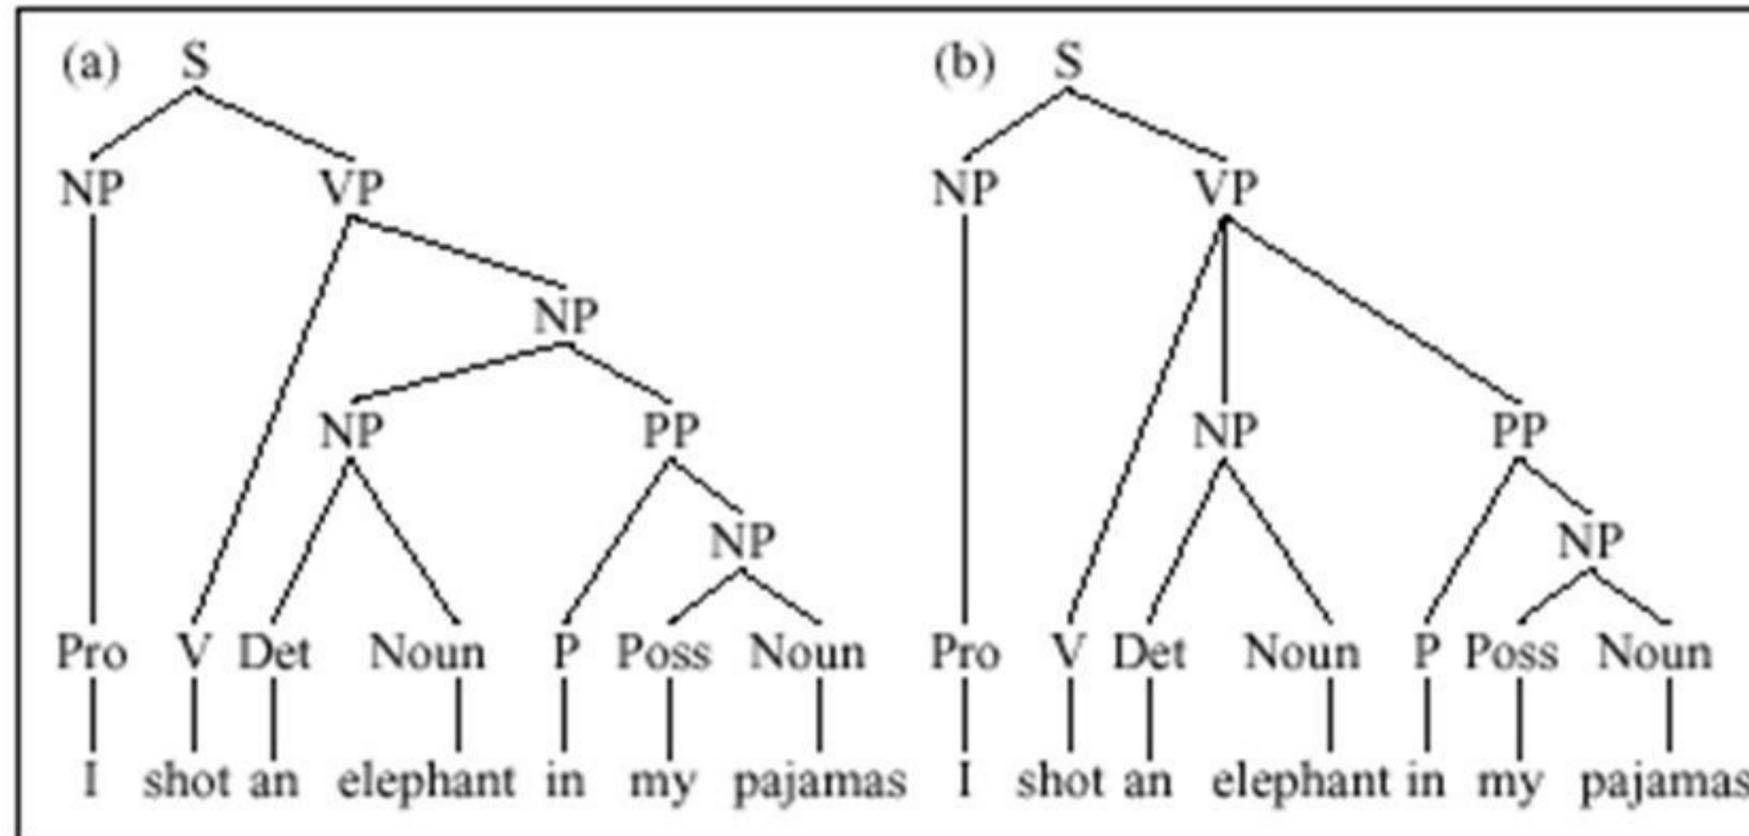
Squad helps dog bite victim

Killer sentenced to die for second time in 10 years

Lack of brains hinders research

# Ambiguity: Syntax

I shot an elephant in my pajamas



# Ambiguity: Semantics

The city council refused the demonstrators a permit because they advocated violence

The city council refused the demonstrators a permit because they feared violence

The city council refused the demonstrators a permit because they \_\_\_\_\_ violence

# Ambiguity?

- Often, there aren't only a few possibilities that can be resolved pragmatically

*il fait vraiment beau* → It is really nice out  
It's really nice  
The weather is beautiful  
It is really beautiful outside  
~~He makes truly beautiful~~  
~~It fact actually handsome~~

- There could be combinatorially many possibilities
- Many won't even register as ambiguities, but systems must resolve them



In the sentence, "The city council refused the demonstrators a permit because they advocated violence" who does "they" refer to ?



In this sentence, the pronoun "they" refers to the demonstrators. The sentence states that



**How did we get so good at natural language processing over the last few years?**

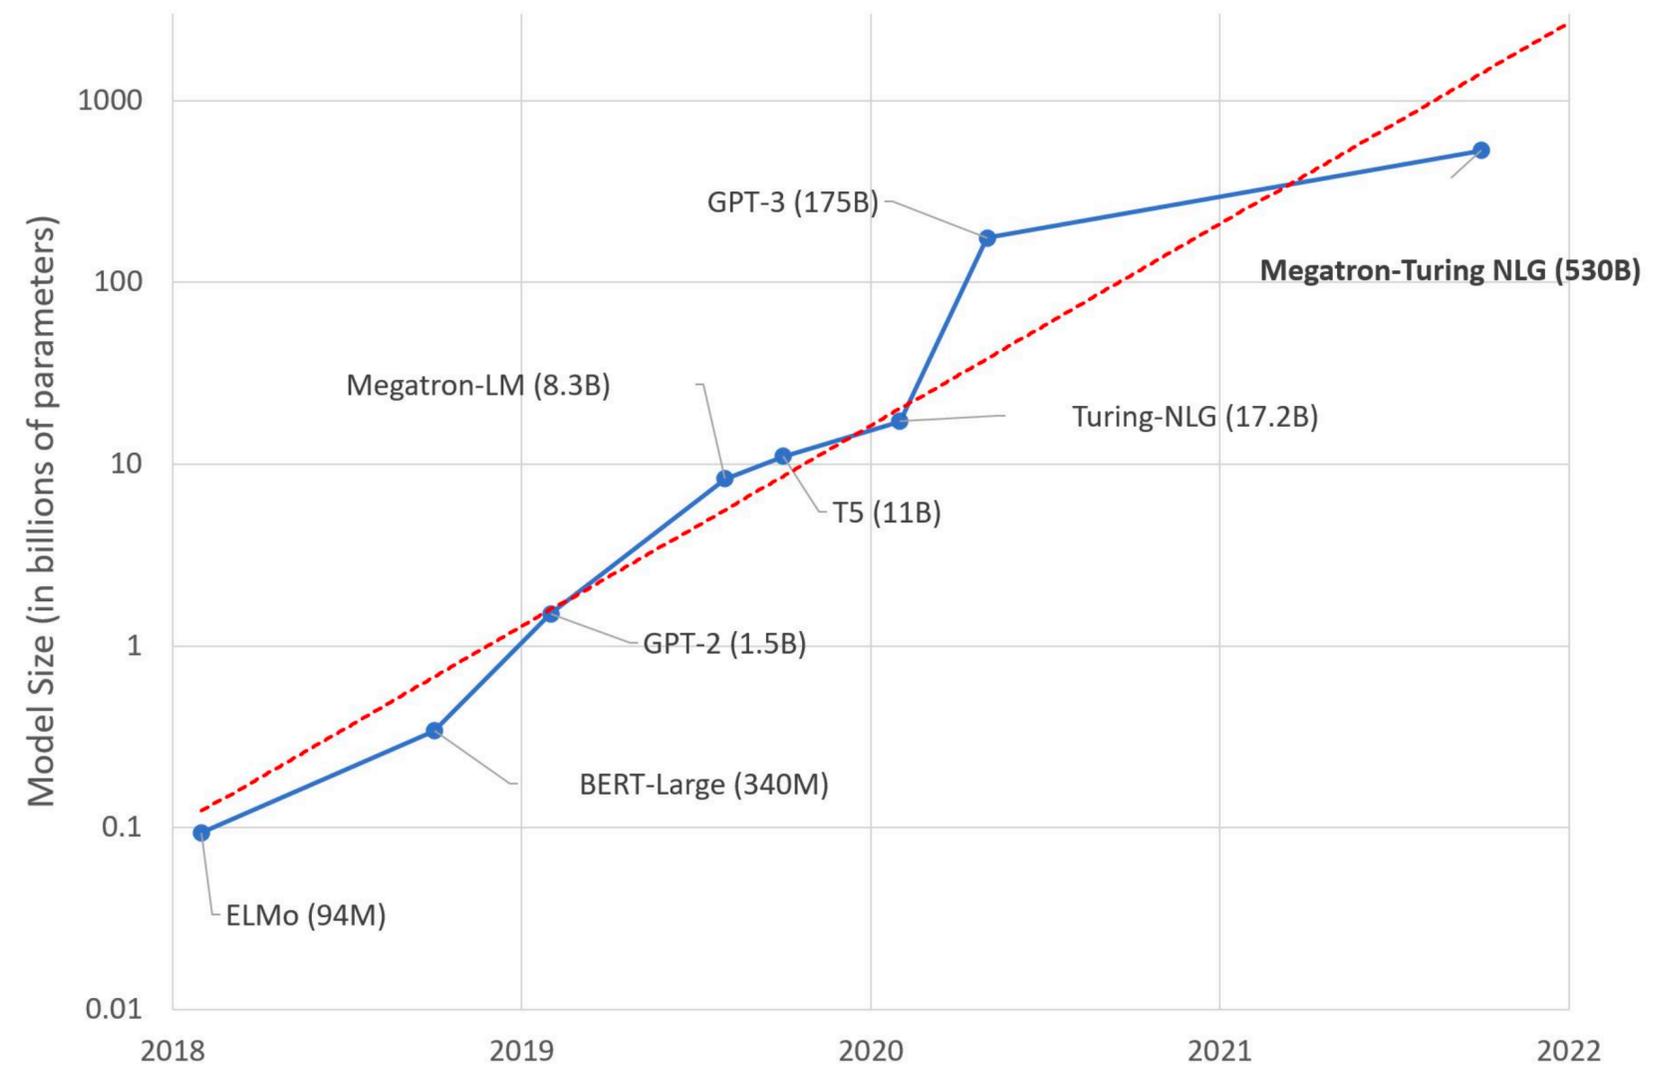


In this sentence, the pronoun "they" refers to the city council. The sentence states that the city council refused the demonstrators a permit because they (the city council) feared violence.



# Language models got a lot bigger!

- Advances in hardware and algorithms allow researchers to train larger language models
- **Growth:** ~5000x increases in model size from 2018 to 2022
- **Comparison:** human cerebral cortex contains ~1000x more neurons than that of mice



AN

In the sentence, "The city council refused the demonstrators a permit because they advocated violence" who does "they" refer to ?



In this sentence, the pronoun "they" refers to the demonstrators. The sentence states that the city council refused the demonstrators a permit because they (the demonstrators) advocated violence.  

**Neural modeling at scale makes things better!**

**Still a lot of work to be done!**



I apologize for the confusion in my previous response. In the sentence, "The city council refused the demonstrators a permit because they feared violence," the pronoun "they" actually refers to the demonstrators.  

The sentence means that the city council denied the permit requested by the demonstrators, and the reason for that denial was that the city council was afraid of violence that might occur if the permit was granted to the demonstrators. So, in this sentence, "they" refers to the demonstrators.

# Today's Outline

- **Section 1:** Course Introduction
- **Section 2:** An simple introduction to Neural NLP

**What will we cover in this class?**

# Course Goals

- Define basic problems and tasks in natural language processing (e.g., text classification; generation)
- Implement common modern approaches for tackling NLP problems and tasks (e.g., embeddings, neural models, transformers)
- Understand failure modes of these models and learning algorithms (e.g., robustness, ethics, evaluation)
- Deploy your understanding of these concepts in applied scenarios (exercise sessions, project)



# Building Blocks: Models

Week	Date	Topic
Week 1	2/18/2026	Introduction + Building a simple neural classifier
	2/19/2026	Neural word embeddings
Week 2	2/25/2026	Classical and Fixed-context Language Models
	2/26/2026	Recurrent Neural Networks + Gated RNN Variants
Week 3	3/4/2026	Attention + Transformers
	3/5/2026	Tokenisation
Week 4	3/11/2026	Pretraining: GPT, BERT, T5
	3/12/2026	<b>Guest Lecture:</b> Science of Pretraining Data

# Building Blocks: Learning

Week 5	3/18/2026	Transfer Learning
	3/19/2026	<b>Guest Lecture:</b> Benchmark Design
Week 6	3/25/2026	Text Generation
	3/26/2026	Text Generation
Week 7	4/1/2026	In-context Learning + Post-training
	4/2/2026	<b>Midterm Review</b>
Week 8	4/8/2026	<b>No Class - Spring Break</b>
	4/9/2026	<b>No Class - Spring Break</b>

# Modern NLP

Week 8	4/15/2026	<b>Midterm</b>
	4/16/2026	<b>Project Description</b>
Week 9	4/22/2026	Interpretability
	4/23/2026	<b>No Class - Work on your project</b>
Week 10	4/29/2026	Ethics in NLP
	4/30/2026	<b>No Class - Work on your project</b>
Week 11	5/6/2026	RAG + Agentic AI
	5/7/2026	<b>Guest Lecture:</b> Frontier AI Agents

# Modern NLP

Week 12 5/13/2026 Scaling Laws + Efficient Deployment

5/14/2026 **Holiday (Ascension)**

Week 13 5/20/2026 Multimodal & Multilingual LLMs

5/21/2026 Wrap-up

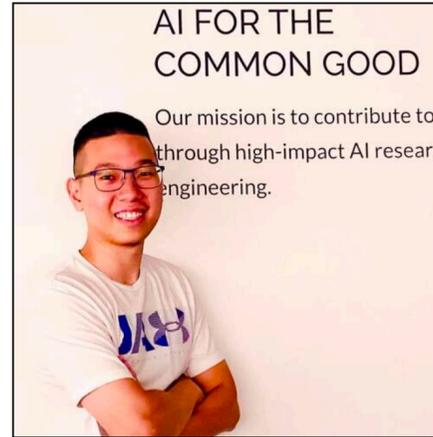
Week 14 5/27/2026 **No Class - Work on your project**

5/28/2026 **No Class - Work on your project**

# Course Team



**Madhur Panwar**  
Head TA



**Zeming (Eric) Chen**  
TA



**Lazar Milikic**  
TA



**Badr AlKhamissi**  
TA



**Sepideh Mamooler**  
TA



**Ayush Tarun**  
TA

# Course Team



**Lysandre Costes**  
AE



**Bich Doan**  
AE



**Petr Grinberg**  
AE



**Sophia Kovalenko**  
AE



**Camille Lannoye**  
AE



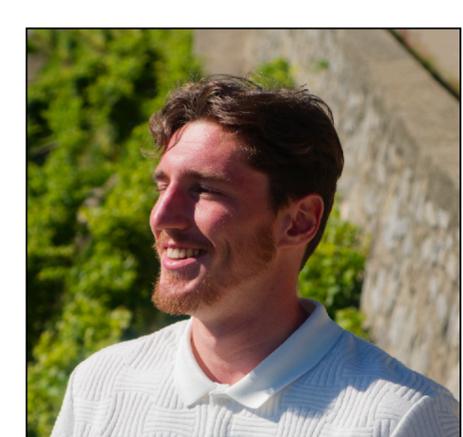
**Sebastian Mayer**  
AE



**Adrian Martínez López**  
AE



**Quentin Sandoz**  
AE



**Emilien Silly**  
AE

# Prerequisites

- This course assumes some prior exposure to the following topics:
  - Comfort with Python Programming (PyTorch, HuggingFace)
  - Comfort with mathematical concepts: linear algebra, probability and statistics
  - Machine Learning (CS-233; CS-433 or some equivalent)
- If you're missing background in these topics, brush up on them, though I'll try keep the course as self-contained!

# Platforms

- Course Webpage: <https://nlp.epfl.ch/cs-552-modern-nlp/>
  - Contains all detailed information about the course (slides, exercise sessions, reading lists, etc.)
- Moodle: <https://moodle.epfl.ch/course/view.php?id=17143>
  - Used to point you to links to other course resources (Webpage, Ed, etc.)
- Ed: <https://edstem.org/eu/courses/3119/discussion>
  - Will be used for class-related questions, discussion, and communication with course staff
  - Queries of a more personal or private nature can be sent to:  
[nlp-cs552-spring2026-ta-team@groupes.epfl.ch](mailto:nlp-cs552-spring2026-ta-team@groupes.epfl.ch)

# Platforms

- **Course Github:** links will be at <https://nlp.epfl.ch/cs-552-modern-nlp/>
  - Contains all detailed information about submitting project milestones & finding exercises

# Lectures

- **Outline Slide:** What will be covered in the class today
- **Motivating Questions:** How should we think about the problems we're trying to solve ?
  - If no one answers, I wait until someone humours me.
- **Content & Material**
- **Comprehension Questions:** Quick tests
  - If no one answers, I go back and re-teach the section :)
- **Participate!**

# Lectures

- Two sessions / week
- **Wednesdays:** 11:15 AM - 1:00 PM in **STCC - Cloud C**
  - **Exceptions:** April 15th (midterm), April 22nd, May 20th. Alternate rooms will be announced.
- **Thursdays:** 1:15 - 2:00 PM in **CE16**
- Lecture and topic schedule found on website
- Lecture slides should be posted before each lecture

# WELCOME TO THE SWISSTECH CONVENTION CENTER

THIS COURSE IS BEING HELD IN A PRIVATE BUILDING. ONLY USE THE ENTRANCE AND FACILITIES  
CLEARLY INDICATED «COURS EPFL»

PLEASE DO NOT EAT OR DRINK IN THE AUDITORIUM – DON'T LEAVE YOUR WASTE

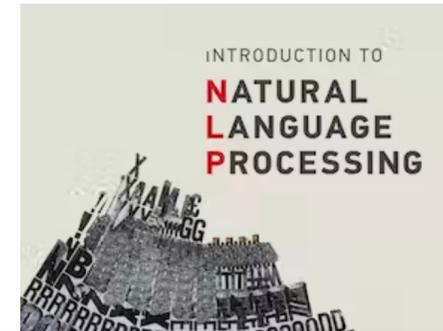
DON'T FORGET TO TAKE ALL YOUR BELONGINGS

SHOULD YOU FORGET SOMETHING, NOTE THAT THE «LOST & FOUND» IS LOCATED AT THE  
ENTRANCE OF THE AUDITORIUM

THANK YOU FOR YOUR COOPERATION !

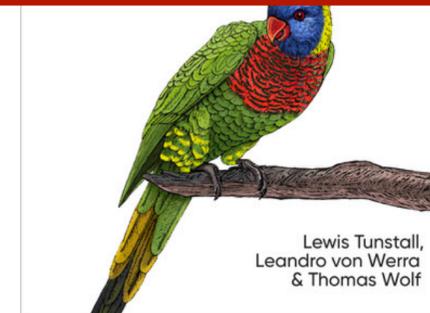
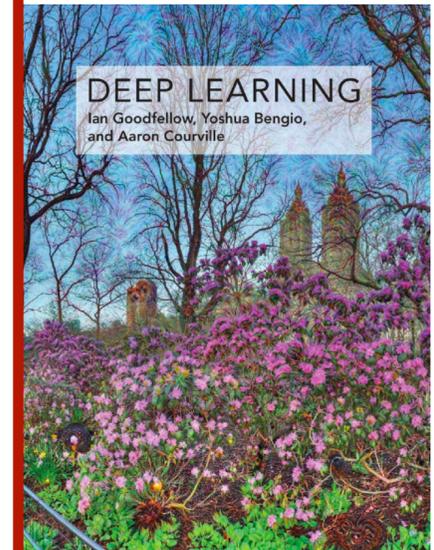
# Textbooks (optional!)

- Jacob Eisenstein. *Natural Language Processing*.
- Yoav Goldberg. *Models for*
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*.
- Lewis Tunstall, Leandro von Werra, and Thomas Wolf. *Natural Language Processing with Transformers*.



**Caution: The NLP field moves very fast!**

Books may be a bit out of date!



# Exercise Sessions

- **Thursdays:** 2:15 - 4:00 pm in **CE11**
- **First half of semester:**
  - Exercises for the week released on Thursdays after last lecture.
  - **Exercise Session Slots:** Review of previous week's exercises
- **Second half of semester:**
  - **Exercise Session Slots:** Optional meetings with Project Supervisors if you have questions about the project
- 1 TA + 2 AEs will be present (could be increased depending on engagement)
- Schedule on website

**How will you be graded in this course?**

# Midterm (30%)

- **90 minutes long**
- Will be held week of **April 15th**
  - Week after Holidays
  - No lecture during week of midterm

# Final Project (70%)

- Complete in teams of 4
- Deliverables (due dates on website):
  - Proposal - Week 10
  - Milestone - Week 12
  - Final report, code, and data - Week 14
- **More details in following weeks!**

# Course Integrity Policy

- For the project, you may build your work upon existing open-source codebases, but are **required to write new code** to perform your experiments. In the project, clearly specify your team's contributions and how they differ from the pre-existing codebase in your milestone reports and final report. **Don't plagiarise.**
- **You can use AI-based tools for any part of your project. Any use of ChatGPT and other AI-based tools must be cited and mentioned. Uncited use of these tools will be penalised.**
- You are free to discuss ideas and implementation details with other project teams. However, you **should not look at another team's code**, or incorporate their code into your project (unless explicitly allowed by the course staff).

# NLP @ EPFL

- **Natural Language Processing** Lab
  - Master's Theses, Semester Projects available every term
- **Swiss AI Initiative: Apertus**
  - Large-scale LLM Training (participation through theses, semester projects, and internships)
- Other **NLP** courses
  - **Fall 2026:** Topics in Natural Language Processing (2 credits)
    - Paper reading, paper reviewing, discussion

**Let's have a great semester!**