

Generative Models on Text Large Language Models

Organization and Overview

Marion Di Marco

April 29, 2025

General Information

- Joint Seminar for Bachelor and Master students
 - first part of the seminar: lectures
 - last 3-4 sessions: student presentations
- Lecturers
 - Professor Alexander Fraser
(Professor for Data Analytics & Statistics)
 - Marion Di Marco (PostDoc)
- Course website:
<https://mariondimarco.github.io/teachingsose25/>

Topics – Outline

- Organization and Introduction to NLP (today!)
- N-Gram models
- Neural LMs and word2vec
- Contextual Word Representations
- Sequence-to-sequence models
- Transformers
- BERT
- RLHF: Training language models to follow instructions with human feedback
- Linguistic information in LLMs
- Student presentations

Student Presentations

- Each student presents a research paper in class
- Research papers from NLP conferences like EMNLP and ACL
 - typically 8-9 pages long
- Topics and groups
 - presentations about related topics per session
- Prepare a presentation of ca. 15-20 minutes
 - present and explain the main ideas of the paper
 - answering questions and discussion in class
- Paper summary

Material and Reading

- *Speech and Language Processing (3rd ed. draft)*,
Dan Jurafsky and James H. Martin (2024)
<https://web.stanford.edu/~jurafsky/slp3/>
- Research papers
- **Preparing for the next lecture:** the reading material for the upcoming lecture will be announced on the course website