

CSC84010/LING8000, Advanced Natural Language Processing
Spring, 2021
Department: Computer Science

Meeting Days, Times, Location: Tuesday 4:15-6:15 pm, CUNY Graduate Center, Online via Zoom

Website: <https://sites.google.com/site/advancednlp21/>

Instructor information:

Prof. Alla Rozovskaya

Email: arozovskaya@qc.cuny.edu

Office location: CUNY Graduate Center, room 5417

Office hours: Tuesday 6:30-7:30 pm

Course Description: This course provides an overview of advanced topics in Natural Language Processing (NLP), with an emphasis on state-of-the-art techniques that require advanced understanding of machine-learning, statistical and deep learning methods that are standard in many NLP applications used today. The course will consist of lectures on advanced topics, paper readings and discussions, and research projects that the students will complete individually or in small groups. The focus of this class will be on developing the necessary skills to read, understand, and implement ideas presented in conference-style NLP papers, as well as to advance the NLP research.

Objectives: This course is intended to provide the students with skills necessary to read and understand research papers published in NLP venues, as well as to develop research projects that will introduce new techniques in the field

Course structure: The course will consist of lectures, paper readings, paper presentations and discussions. The main component of the course is the research project.

Readings: The required textbook for the course is Jurafsky and Martin, [*Speech and Language Processing, 3rd edition*](#). Readings will also be assigned from published notes and tutorials that will be made available online.

Background and prerequisites: Machine Learning and/or Introduction to Natural Language Processing. Programming experience is necessary to complete the research project.

Textbook: D. Jurafsky and J. H. Martin [Speech and Language Processing, 3rd ed.](#)

<https://web.stanford.edu/~jurafsky/slp3/>

Topics:

- Neural networks: feed-forward neural networks, recurrent neural networks, transformers
- Byte-pair encoding, neural machine translation
- Linear learning algorithms: Naïve Bayes, Logistic Regression
- Vector Space Models
- Word embeddings
- BERT, RoBERTa
- NLP applications in the context of the above techniques

Assessment/Grading

- Paper presentations (30%)
- Class participation (20%).
- Final project 50%)

Policy:

Absentees are solely responsible for catching-up. Academic dishonesty, such as plagiarism or cheating - taking other people's work with or without their permission in order to get credit for yourself, will be dealt with seriously, including an "F" grade for the course and/or disciplinary action according to the University's policy on academic integrity.

Honor Code applies: The work you turn in must be your own.

Note: If any student needs special accommodations because of a disability, please contact the instructor during the first week of classes.

CUNY Policy on Academic Integrity: Academic Dishonesty is prohibited in The City University of New York and is punishable by penalties, including failing grades, suspension, and expulsion as provided at: <http://web.cuny.edu/academics/info-central/policies/academic-integrity.pdf>. Please read this document, paying careful attention to the sections on plagiarism and Internet plagiarism. If you are not sure how to cite work you have found on the internet, please review the APA Guidelines provided by the [Purdue OWL](#).