

CS 4984: Natural Language Processing

Fall 2015

Dr. Alla Rozovskaya

Course description: This course provides an introduction to natural language processing (NLP), a field concerned with automatic processing of “natural” (human) language. The course will study modern computational approaches used in the processing of fundamental linguistic phenomena. The emphasis will be on machine learning and statistical methods. The materials will include basic topics such as language modeling, part-of-speech-tagging, parsing, as well as various natural language processing applications, such as machine translation.

Objectives: This course is intended to provide students with general understanding of the key linguistic concepts relevant to the processing of natural language, to help students acquire knowledge of the state-of-the-art computational techniques used in different areas of natural language processing, and to develop basic skills required for reading and understanding research papers in NLP.

Course structure: the course will consist of lectures, homework assignments that will include programming exercises, a midterm exam and a final exam.

Textbook: Jurafsky and Martin (2008), [Speech and Language Processing, 2nd edition](#). This is a required text.

Target audience and prerequisites: Advanced undergraduate students and graduate students. Programming experience is necessary for the assignments (CS 3114). Prior exposure to statistics (e.g., STAT 4705 or 4714) and algorithms (e.g., CS 4104) is strongly recommended. Linguistic background is not required.