

CS 6120/CS 4120: Natural Language Processing

Professor Lu Wang

Spring 2019

E-mail: luwang@ccs.neu.edu

Office Hours: F 3pm-4pm (or by appointment)

Office: 911, 177 Huntington Ave.

Web: <http://www.ccs.neu.edu/home/luwang/>

Class Hours: TF 9:50am -11:30am

Class Room: West Village H 110

Course Description

This course aims to introduce fundamental tasks in natural language processing, and its recent advances based on machine learning algorithms (e.g., neural networks) and applications for interdisciplinary subjects (e.g., computational social science). The course materials are mostly delivered as lectures, and accompanied with reading materials. The students will be evaluated based on assignments, a research-driven course project, and an open-book final exam.

Course and Discussion Websites

Course Website: http://www.ccs.neu.edu/home/luwang/courses/cs6120_sp2019/cs6120_sp2019.html

Discussion Forum: <http://piazza.com/northeastern/sp2019/cs6120/home>

Required Materials

- Main textbooks
 - Dan Jurafsky and James H. Martin, "Speech and Language Processing, 2nd Edition", Prentice Hall, 2009.
 - Third edition draft is available at web.stanford.edu/~jurafsky/slp3/.
- Other Reference
 - Chris Manning and Hinrich Schütze, "Foundations of Statistical Natural Language Processing", MIT Press, 1999

- Since many natural language processing problems are driven by machine learning techniques nowadays, we also highly encourage you to read machine learning textbooks:
 - Christopher M. Bishop, “Pattern Recognition and Machine Learning”, Springer, 2006.
 - Tom Mitchell, “Machine Learning”, McGraw Hill, 1997.

Prerequisites/Corequisites

This course is designed for graduate students and senior undergraduate students majoring in computer science, linguistics, and other related areas. Students who take this course are expected to be able to write code in some programming languages (e.g., Python, Java, or C/C++) proficiently, and finish courses in algorithms (CS 5800 or CS 7800), multivariable calculus, probability, and statistics. Linear algebra is optional, but highly recommended. It would be beneficial if the students have prior knowledge on supervised machine learning.

Course Structure

Grading Policy

Blackboard is used for electronic submission. Hard copies are submitted in class. Assignment or report turned in late will be charged 20 points (out of 100 points) off for each late day (i.e. 24 hours). Each student has a budget of 5 days throughout the semester before a late penalty is applied. You may want to use it wisely, e.g. save for emergencies. Each 24 hours or part thereof that a submission is late uses up one full late day. Late days are not applicable to final presentation. Each group member is charged with the same number of late days, if any, for their submission.

Grades will be determined based on two assignments, eight in-class tests, one course project, one open-book exam, and participation:

- 30% of your grade will be determined by two assignments, each of 15%.
- 5% of your grade will be determined by eight quick in-class tests, each of 1%; three with lowest grades will be dropped, and no make-up.
- 25% of your grade will be determined by your course project, including proposal (3%), progress report (7%), final report (10%), and presentation (5%) with 1% as bonus if selected as best project by peer student.
- 35% of your grade will be determined by final exam.
- 5% of your grade will be determined by your participation in the class.

Course Policies

Academic Integrity and Honesty

This course follows the Northeastern University Academic Integrity Policy. All students in this course are expected to abide by the Academic Integrity Policy. Any work submitted by a student

in this course for academic credit should be the student's own work. Collaborations are allowed only if explicitly permitted. Violations of the rules (e.g. cheating, fabrication, plagiarism) will be reported.

Schedule and Weekly Learning Goals

The schedule is tentative and subject to change. The learning goals below should be viewed as the key concepts you should grasp after each week.

Week 01, 01/07 - 01/11: Topic: Introduction, Language Models

- Reading: Ch1, Ch4.1-4.9
- TODO: start thinking about projects and looking for teammates

Week 02, 01/14 - 01/18: Topic: LM cont'd, Text Categorization and Evaluation, Naive Bayes, Part-of-Speech Tagging, Sequence Labeling, Hidden Markov Models

- Reading: Ch5.1-5.5, Ch6.1-6.5

Week 03, 01/21 - 01/25: Topic: HMM cont'd, Word Sense Disambiguation

- Reading: Ch20

Week 04, 01/28 - 02/01: Topic: Machine Learning Basics (Maximum Entropy, Feedforward Neural Networks), Formal Grammars of English, Syntactic Parsing

- Reading: Ch12.1-12.5

Week 05, 02/04 - 02/08: Topic: Parsing cont'd, Dependency Parsing, Semantics

- Reading: Ch12.1-12.5, Ch13.1-13.4.1

Week 06, 02/11 - 02/15: Topic: Vector-Space Lexical Semantics, Semantic Parsing, Combining Logical and Distributional Semantics Information Extraction

- Reading: Ch20.7

Week 07, 02/18 - 02/22: Topic: Information Extraction, Neural Networks

- Reading: Ch22.1-22.2

Week 08, 02/25 - 03/01: Neural Language Models, Sentiment Analysis, Opinion Mining, NLP and Social Media

- NO CLASS is planned on 03/01
- Reading: Ch23

Week 09, 03/04 - 03/08: Topic: (NO CLASS: spring break)

Week 10, 03/11 - 03/15: Topic: Text Summarization, Discourse Analysis

- Reading: Ch23.3-23.7, Ch21

Week 11, 03/18 - 03/22: Topic: Project progress feedback, Question Answering

- Reading: Ch23.1-23.2

Week 12, 03/25 - 03/29: Topic: Dialogue Systems and Chatbots

- Reading: Ch24

Week 13, 04/01 - 04/05: Topic: Machine Translation

- Reading: Ch25.1-25.5

Week 14, 04/08 - 04/12: Topic: Project Presentation

- TODO: Project final report due on December 10.

Week 15, 04/15 - 04/19: Topic: Final Exam

- Pending confirmation.