

Syllabus

Introduction to Natural Language Processing Spring 2017 CAP 4641: Undergraduate / CAP 5640: Graduate

Taught by

Professor Mark A. Finlayson

ECS Room 362

markaf@fiu.edu / (305) 348-7988 (office) / markaf33 (skype) / markaf@fiu.edu (Google+)

Summary

This class is an introduction to fundamental principles and techniques of natural language processing (NLP). In this class you will become familiar with core NLP principles, techniques, and algorithms. In homeworks and exams you will tackle key questions and problems in the NLP field, as well as simulate step-by-step a number of fundamental NLP algorithms. Both the graduate and undergraduate version of the course share the readings, lectures, homeworks, and exams. Graduate students additionally will carry out an independent project, the goal of which is to understand in depth a specific NLP task or problem through detailed reading and hands-on implementation. The ultimate goal of this course is to set the stage for students to engage in research that advances the state of the art of NLP.

Textbook

Daniel Jurafsky & James H. Martin, *Speech and Language Processing*, 2nd Edition, Prentice Hall, ISBN-13: 978-0131873216

Times & Locations

Lecture will be held on Tuesdays and Thursdays from 5pm to 6:15pm. Exceptions are noted on the website.

Office hours will be held in ECS 362 on Mondays from 3:30pm to 5pm.

Remote students should call my office during office hours, or contact me via email to arrange a skype or Google+ call. I am available outside office hours by appointment.

Grading

Item	Undergraduate (CAP 4641)	Graduate (CAP 5640)
Homework (10)	32%	24%
Midterm Exam	32%	24%
Final Exam	32%	24%
Project	-	24%
Participation (attendance will be noted)	4%	4%
Total	100%	100%

Reading

Reading assignments will be distributed via the website. Readings will be associated with each lecture, and these should be completed before lecture begins, as we will rely on content in the reading during lecture.

Homework

Version of November 2, 2016, 9:17 AM

There will be 10 homework assignments, appearing approximately every week during the term (with a few exceptions—see the website). Homeworks will be available on the website by the beginning of lecture on Tuesday each week, and will be due the following Tuesday at 4pm. Homeworks must be submitted electronically as a single pdf file via the website; homeworks submitted in other formats will be returned. If you need an extension due to extenuating circumstances, please don't hesitate to ask.

Exams

There will be two exams: a midterm and a final. The midterm will be held on February 25, during normal class time. You will be allowed one single letter-sized page of notes (double-sided, may be printed if you wish). The final (cumulative) will be held on Tuesday, May 3rd from 5 to 7pm, during which you will be allowed two letter-sized pages of notes (double-sided).

Project

The class project (only for graduate students) will proceed in five stages. Project parts are due by 5pm on the indicated Fridays.

1. Proposal (10%; due end of week 2)
2 pages outlining an NLP task of interest to the student, with proposed seed references and final implementation deliverables. This will set the stage for the literature review. Students are encouraged to come speak with the professor about their project ideas.
2. Literature Review (20%; due end of week 6)
10 pages (minimum) reviewing the literature relevant to the proposed project. This review will give the student the necessary background to proceed to project implementation.
3. Summary (5%; due end of week 8)
1 paragraph summarizing project deliverables as modified to reflect feedback given for the literature review. At this stage students may wish to assemble into teams, and the purpose of the summary is to allow merging of individual student's project work so far. Students are encouraged to come meet with the professor to discuss merging. Each team (or individual student, if working alone) should prepare a 1 paragraph summary clearly outlining what exactly the project will entail, and laying out clear and unambiguous criteria for success.
4. Implementation & Report (50%; due last week of class)
20 pages (minimum) describing the project implementation, what was done, what was achieved, and how it was measured. Portions of the student's literature review should be incorporated as appropriate.
5. Presentation (15%; in lecture, last week of class)
A presentation to the class on the project, its goals, and results. Presentation length will be determined by the total number of teams, divided equally across the 2 final classes.

Academic Integrity

Please review the discussion of academic integrity in the FIU student handbook: we will adhere to those definitions and procedures. You are welcome and encouraged to collaborate with your fellow students on the homeworks and when studying for exams; however, you must write up your own solutions, in your own words. Directly copying solutions will be considered plagiarism and treated accordingly. If you do collaborate on homeworks, please note your collaborators on the first page.