

USC Viterbi School of Engineering

CSCI 544: Applied Natural Language Processing

Units: 4

Term—Day—Time:

Spring 2022 – Mon/Wed – 3:00-4:50pm

Location: SGM 123

Instructor: Xuezhe Ma

Office Hours: After each class virtually, or by appointment

Contact Info: xuezhema@isi.edu

Instructor: Mohammad Rostami

Office Hours: Mon 1-2 PM and virtually at:

<https://usc.zoom.us/my/mrostami>

Contact Info: mrostami@isi.edu

Teaching Assistants:

Office Hours: Tue 11-12 AM and virtually at:

Contact Info:

Office Hours:

Contact Info:

Office Hours:

Contact Info:

Office Hours:

Contact Info:

Graders:

Catalogue Course Description

This course covers both fundamental and cutting-edge topics in Natural Language Processing (NLP) and provides students with hands-on experience in NLP applications.

Learning Objectives

The syllabus is available at the following link:

https://docs.google.com/document/d/1gogX7IAqHRdWlgsO_EMtPhUghFnIG6Ux/

PLEASE CHECK THE LINK CONTINUALLY FOR THE LATEST VERSION

- Read technical literature in Natural Language Processing (including original research articles) and answer questions about such readings.
- Implement language processing algorithms and test them on natural language data.
- Solve language processing problems and explain the reasoning behind their solution

Required Preparation:

Experience programming in Python

Course Notes

The course will be run as a lecture class with student participation strongly encouraged. There are weekly readings and students are encouraged to do the readings prior to the discussion in class. All of the course materials, including the readings, lecture slides, and homeworks will be posted online. The class project is a significant aspect of this course and at the end of the semester students will present their projects in the form of short videos.

Required Readings and Supplementary Materials

Textbook:

Foundations of Statistical Natural Language Processing by Manning and Schütze

Speech and Language Processing by Jurafsky and Martin (3rd edition draft),

We use a set of technical papers and book chapters that are all available online. All of the required readings are listed in the course schedule.

Description and Assessment of Assignments

Homework Assignments

There will be four coding homework assignments. The assignments must be done individually. Each assignment is graded on a scale of 0-10 and the specific rubric for each assignment is given in the assignment.

Grading inquiries and questions about the grading of the homeworks and the quizzes can be asked (to the TA) within two weeks from the grading date.

Course Project

An integral part of this course is the course project, which builds on the topics and techniques covered in the class. Students can work in teams of five people on their project.

Project Timeline:

- Week 6: Project proposals (2 pages)
- Week 10: Project status update due (one page status report)
- Week 13: Project final report (4 pages) and short videos (2 minutes)

Project description: Each project team will select a topic of their choice. The project types can include NLP prototype design, presenting the design of a novel, original NLP application.

Grading breakdown of the course project:

- Proposal: 10%
- Status Reports: 10%
- Project Presentation: 10%
- Final Write-up: 70%

Grading Breakdown

Quizzes: There will be weekly quizzes at the start of class based on the material from the week before. The **highest ten quiz grades** will be considered. Missed quizzes will receive a zero grade, and there will be no make-up quizzes for any reason.

Homework: There will be four coding homework based on the topics of the class.

Final Exam: There is a multiple choice final exam at the end of the semester covering all of the material covered in the class. The final exam will be held on the date designated by USC

Class Project: Each student will do a group class project based on the topics covered in the class. Students will propose their own project, do the research and build a proof-of-concept, create a video demonstration of the proof-of-concept, and present the project in their report.

Paper Presentation: The research article presentation is an activity where project teams read an article of their choice and present it to the class. It does not have to be related to the team's project. You will read the article, identify the central points of the research, and present that research to the class. You will also comment on research article presentations of fellow students.

Grading Schema:

Quizzes	10%
Homework	40%
Paper Presentation:	10%
Class Project	30%
Final	10%
<hr/>	
Total	100%

Grades will range from A through F. The following is the breakdown for grading:

	74 – 76.9 = C+	Below 60 is an F
90 – 100 = A	70 – 73.9 = C	
87 – 89.9 = A-	67 – 69.9 = C-	
84 – 86.9 = B+	64 – 66.9 = D+	
80 – 83.9 = B	62 – 63.9 = D	
77 – 79.9 = B-	60 – 61.9 = D-	

Assignment Submission Policy

Homework assignments are due at 2:59pm on the due date and should be submitted on Blackboard. Late submissions without prior approval, e.g., due to medical conditions, will not be submitted.

Course Schedule: A Weekly Breakdown

#	Date	Lecture	Reading	Instructor
1	01/09/2022	Introduction	Jurafsky and Martin, Speech and Language Processing (3rd edition draft), Chapter 2: Regular Expressions, Text Normalization, and Edit Distance.	XM
2	01/11/2022	Naive Bayes, Linear Classifier	Jurafsky and Martin, Speech and Language Processing (3rd edition draft), Chapter 4: Naive Bayes Classification and Sentiment Analysis HW1 Release	XM
3	01/18/2022	Perceptron, SVM, Logistic Regression,	Jurafsky and Martin, Speech and Language Processing (3rd edition draft), Chapter 5: Logistic Regression.	MR

		Model evaluation	Hal Daumé III, A Course in Machine Learning (v. 0.99 draft), Chapter 4: The Perceptron. Tutorial on Support Vector Machine (SVM) by Vikramaditya Jakkula	
4	01/23/2022	Sequence Labeling & HMMs	Jurafsky and Martin, Speech and Language Processing (3rd edition draft), Chapter 8: Part-of-Speech Tagging Jurafsky and Martin, Speech and Language Processing (3rd edition draft), Appendix A: Hidden Markov Models Michael Collins: Tagging Problems, and Hidden Markov Models	MR
5	01/25/2022	Discriminative and Unsupervised Sequence Labeling	Brill (1992): A Simple Rule-Based Part of Speech Tagger Ratnaparkhi (1996): A Maximum Entropy Model for Part-Of-Speech Tagging Lafferty, McCallum and Pereira (2001): Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data Michael Collins: Log-Linear Models, MEMMs, and CRFs HW1 Deadline HW2 Release	MR
6	01/30/2022	Word Embedding	Mikolov, Tomas, et al. "Efficient estimation of word representations in vector space." arXiv preprint arXiv:1301.3781 (2013). Mikolov, Yih and Zweig (2013): Linguistic Regularities in Continuous Space Word Representations Jurafsky and Martin, Speech and Language Processing (3rd edition draft), Chapter 6	MR
7	02/01/2022	Word Embedding (continued), Neural Network	Jurafsky and Martin, Speech and Language Processing (3rd edition draft), Chapter 7	MR
8	02/06/2022	PyTorch & Basic Concepts in DL		MR

9	02/08/2022	Language Models, smoothing	Jurafsky and Martin, Speech and Language Processing (3rd edition draft), Chapter 3: N-gram Language Models. HW2 Deadline HW3 Release	MR
10	02/13/2022	Spelling Correction & Statistical Machine Translation	Notes by Michael Collins Spelling Correction and the Noisy Channel	MR
11	02/15/2022	Expectation Maximization for MT	Michael Collins, The Naive Bayes Model, Maximum-Likelihood Estimation, and the EM Algorithm	MR
12	02/22/2022	Phrase Based Translation, MT Decoding	Statistical Phrase-Based Translation, Koehn et al Project Proposal Deadline	MR
13	02/27/2022	Dialogue	Levin, Pieraccini and Eckert (2000): A Stochastic Model of Human-Machine Interaction for Learning Dialog Strategies Jurafsky and Martin, Speech and Language Processing (3rd edition draft), Chapter 24: Chatbots and Dialog Systems	MR
14	03/01/2022	Deep Learning for NLP	HW3 Deadline HW4 Release	XM
15	03/06/2022	RNN & Text classification		XM
16	03/08/2022	Sequence-to-sequence Learning & Machine Translation		XM
17	03/20/2022	Transformers-I	Attention is All You Need HW4 Deadline	XM
18	03/22/2022	Transformers-II		XM
19	03/27/2022	Advanced Topics in Machine Translation	Status Report Deadline	XM

20	03/29/2022	Neural Language Models & Contextualized Embeddings-I		XM
21	04/03/2022	Neural Language Models & Contextualized Embeddings-II	Paper Selection Deadline	XM
22	04/05/2022	Deep Generative Models		XM
23	04/10/2022	Paper Presentation		
24	04/12/2022	Paper Presentation		
25	04/17/2022	Paper Presentation		
26	04/19/2022	Paper Presentation		
27	04/24/2022	Project Presentation		
28	04/26/2022	Project Presentation	Project Final Report	

Statement on Academic Conduct and Support Systems

Academic Conduct

Plagiarism – presenting someone else’s ideas as your own, either verbatim or recast in your own words – is a serious academic offense with serious consequences. Please familiarize yourself with the discussion of plagiarism in *SCampus* in Section 11, Behavior Violating University Standards <https://scampus.usc.edu/1100-behavior-violating-university-standards-and-appropriate-sanctions>. Other forms of academic dishonesty are equally unacceptable. See additional information in *SCampus* and university policies on scientific misconduct, <http://policy.usc.edu/scientific-misconduct>.

Discrimination, sexual assault, and harassment are not tolerated by the university. You are encouraged to report any incidents to the *Office of Equity and Diversity* <http://equity.usc.edu> or to

the *Department of Public Safety*

<http://capsnet.usc.edu/department/department-public-safety/online-forms/contact-us>. This is important for the safety of the whole USC community. Another member of the university community – such as a friend, classmate, advisor, or faculty member – can help initiate the report, or can initiate the report on behalf of another person. *The Center for Women and Men* <http://www.usc.edu/student-affairs/cwm/> provides 24/7 confidential support, and the sexual assault resource center webpage <http://sarc.usc.edu> describes reporting options and other resources.

Support Systems

A number of USC's schools provide support for students who need help with scholarly writing. Check with your advisor or program staff to find out more. Students whose primary language is not English should check with the *American Language Institute* <http://dornsife.usc.edu/ali>, which sponsors courses and workshops specifically for international graduate students. *The Office of Disability Services and Programs* http://sait.usc.edu/academicsupport/centerprograms/dsp/home_index.html provides certification for students with disabilities and helps arrange the relevant accommodations. If an officially declared emergency makes travel to campus infeasible, *USC Emergency Information* <http://emergency.usc.edu> will provide safety and other updates, including ways in which instruction will be continued by means of blackboard, teleconferencing, and other technology.