Computational Linguistics (LING 28610/38610, CMSC 25610/35610)

Fall 2019

Course Information
Computational Linguistics (LING 28610/38610, CMSC 25610/35610)
Times: Tuesday and Thursday 11:00am–12:20pm
Location: Oriental Institute 208

Contact Information
Instructor: Allyson Ettinger
Office: Rosenwald 229B
Email: aettinger@uchicago.edu
Office hours: Mon 2:00-3:00pm (or by appointment)

Other information
Course website: `https://canvas.uchicago.edu/courses/23821`
Prerequisite: Instructor permission. Students should have two courses in linguistics or other cognitive science areas, or equivalent background in relevant topics.
Text: There is no textbook for this class. All readings will be provided on Canvas.

Course description: "Computational linguistics" describes a range of research areas that use computational approaches to address scientific and engineering problems pertaining to language. The goal of this course is to give students a foundational understanding of two domains of computational linguistics: cognitive modeling and natural language processing. Cognitive modeling uses computational models to test scientific hypotheses about mechanisms underlying human cognition. Natural language processing designs computational models of language for engineering purposes, in service of advancement of artificial intelligence. In this course we will discuss various types of computational models and their applications within each of these subdomains.

On the technical side, the additional goal of this course is to afford students the basic programming skills and mathematical foundations to be able to implement computational models and conduct research in computational linguistics. This course will thus combine reading and discussion of relevant research with instruction and practice in programming and mathematical foundations.

Expectations and grading procedures:

1. **Participation.** This course will combine lecture-style instruction with full-class discussion, so you will be expected to do readings and participate regularly in class discussions.

2. **Exercises and problem sets.** There will be regular programming exercises and problem sets aimed at implementing ideas discussed in class, while also bringing everyone up to speed with programming capabilities. These assignments will use Python.

3. **Final project.** You will complete a project implementing a computational model to address a scientific or engineering problem, or executing a substantive analysis of an existing computational model. Required components of this project are as follows:

   Final project components

   Proposal. Due Oct 31. By this time you need to have met with me to discuss your choice of project and obtain approval.

   Draft. Due Nov 12. Submit a rough introduction and literature review for your paper. The introduction does not need to state your results (which you may not have in full) but should clearly frame the motivation and intended contribution of the work.

   Presentation. In class Nov 26 and Dec 3. You will give a short presentation of your project in class. You should bring a handout and/or use slides. Further guidelines will be announced.

   Paper. Due Wednesday Dec 11. Submit online. Further guidelines will be announced.

Grading will be weighted as follows:

Participation: 15%
Programming exercises / problem sets: 40%
Final project: 45% (Draft, 10%, Presentation 5%, Paper 30%)

Initial course schedule (subject to change!):
Any changes to this schedule will be announced via Canvas.

Date	**Topic**	**Reading**	**Assignment due**
		Intro	
Tuesday 10/1	Course introduction and overview		
		Word Learning	
Thursday 10/3	Modeling word learning dynamics; probability intro	McMurray (2007); Russell & Norvig (2010) Ch. 13	Programming exercise: self intro
Tuesday 10/8	Bayesian word learning	Xu & Tenenbaum (2007)	Programming exercise: McMurray
Thursday 10/10	Vector space representations for words	Turney & Pantel (2010); Landauer & Dumais (1997)	Programming exercise: Bayes
		Machine learning intro (supervised)	
Tuesday 10/15	Classification intro	*Course in Machine Learning* Ch. 1; Jurafsky & Martin Ch. 4	Programming exercise: vector space models
Thursday 10/17	Neural network intro	Jurafsky & Martin Ch. 7	Programming exercise: classifiers

Words in context			
Tuesday 10/22	Language modeling, information theory, and connections to psycholinguistics	Jurafsky & Martin Ch. 3	Programming exercise: neural networks
Thursday 10/24	Neural network language models	Bengio et al. (2003)	
Bringing in syntax			
Tuesday 10/29	Constituency grammar, parsing	Jurafsky & Martin Ch. 11; Ch. 12	Language modeling problem set
Thursday 10/31	Syntactic surprisal	Hale (2001)	Deadline to have met to discuss project
Tuesday 11/5	Noisy channel models and application to processing	Levy (2008)	
Contextualizing word representations			
Thursday 11/7	ELMo: adding context information to word representations	Peters et al. (2018)	Psycholing. modeling problem set
Tuesday 11/12	BERT: pushing the state of the art with contextualized representations and transformer networks	Devlin et al. (2018)	Due: draft intro / literature review
Evaluation/analysis			
Thursday 11/14	Neural decoding	Mitchell et al. (2008)	
Tuesday 11/19	Probing tasks	Adi et al. (2016); Conneau et al. (2017); Ettinger et al. (2018)	
Thursday 11/21	Biases in datasets	Kaushik & Lipton (2018)	Analysis problem set
Presentations			
Tuesday 11/26	Project presentations		
Thursday 11/28	*Thanksgiving break*		
Tuesday 12/3	Project presentations		
Finals			
Wednesday 12/11			Final paper due