

# Computational Linguistics Seminar (LING 58600)

Spring 2020

## Course Information

*Computational Linguistics Seminar* (LING 58600)

Time: Wednesday 10:30am–1:20pm

Location: Zoom!

## Contact Information

Instructor: Allyson Ettinger

Office: See Canvas for meeting Zoom link

Email: aettinger@uchicago.edu

Office hours: Weds 2:00-3:00pm, or by appointment

## Other information

Course website: <https://canvas.uchicago.edu/courses/27700>

Prerequisite: Preferably some background in areas of cognitive science and/or computer science.

Text: There is no textbook for this class. All readings will be provided on Canvas.

**Course description:** Perhaps the most fundamental contribution of human language is the capacity to convey meaning. In this course we will examine approaches to meaning in computational linguistics—particularly, efforts to capture meaning and “understanding” in artificial intelligence. Recreating linguistic meaning in artificial intelligence is fraught with challenges: not only the engineering challenges of simulating human language comprehension, but also the challenge of defining what it means to comprehend language, and of assessing the extraction and representation of meaning in complex black box systems used in modern artificial intelligence. In exploring these problems, we stand not only to improve the performance of artificial intelligence systems, but also to gain insight into how semantic processing and language comprehension operate in the human mind. In this seminar we will discuss a range of literature tackling the problem of extracting, processing, and representing meaning in machines, and we will explore in detail how these approaches relate to the functioning of linguistic meaning in humans. We will also explore current standards by which “language understanding” is evaluated in artificial intelligence, and critique these methods from the perspective of human language comprehension and semantic processing.

## Expectations and grading procedures:

1. **Participation.** This course is a discussion-based seminar, so it will be critical that you do the readings and participate regularly in class discussions.
2. **Discussion posts.** You will also be expected to post on Canvas a thoughtful, substantive comment and/or question about the reading, before the start of the class period for which that reading is assigned, for each class period. Because the course will be held via Zoom, and

three hours is a long time even in person, we will likely shorten the synchronous discussion component. If we do so, the discussion posts will be an even more critical component of the class than usual. This will involve posting to the Canvas discussion board a thoughtful, substantive comment and/or question about the reading, before the start of the class period for which that reading is assigned, for each class period.

3. **Final project.** You will complete a project pertaining to some aspect of meaning in computational systems. I will be available throughout the quarter to discuss, and help in formulating, project ideas. Detailed guidelines for the final project will be made available later in the quarter, but you will be able to choose between a variety of project types, including the following:

- Propose a model to address an aspect of meaning in NLP.
- Design an evaluation for the quality of meaning representations or “understanding” in computational systems.
- Do a literature review on a topic relevant to meaning in NLP.

### Final project components

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

*Paper.* Due June 11. Submit online. Further guidelines will be announced.

### Grading will be weighted as follows:

In-class participation and reading posts: 50%

Final project: 50% (Presentation 5%, Paper 45%)

### Course schedule

(subject to change—check Canvas for the most updated syllabus and assignments):

Date	Topic	Reading	Assignment due
<b>Intro</b>			
Wednesday 4/8	Course introduction, hand-designed databases for meaning	Stringer (2019), Miller et al. (1993)	Discussion post
Wednesday 4/15	Vector space models for representing word meaning	Landauer and Dumais (1997), Mikolov et al (2013)	Discussion post
Wednesday 4/22	Tuning vector space models with structured linguistic information	Faruqui et al. (2015), Jauhar et al. (2015), Levy et al. (2014)	Discussion post
Wednesday 4/29	Compositional semantics and composition of word vectors	Heim and Kratzer (1998) excerpts, Baroni & Zamparelli (2010), Kintsch (2001)	Discussion post
Wednesday 5/6	Annotating and predicting components of sentence meaning	Marquez et al. (2008), Banarescu et al. (2013)	Discussion post

Wednesday 5/13	Generalized sentence encoders	Socher et al. (2012), Bowman et al. (2016), Conneau et al. (2017)	Discussion post
Wednesday 5/20	Recent developments in “natural language understanding” engineering and evaluation	Devlin et al. (2018), Wang et al. (2018)	Discussion post
Wednesday 5/27	Meaning meets world: learning and testing of common-sense reasoning	Zellers et al. (2018), Talmor et al. (2019)	Discussion post
Wednesday 6/3	Final project presentations and wrap-up discussion		Project presentation
Wednesday 6/11			Final project paper due