Introduction to Computational Linguistics and Large Language Models

Instructor: Hayley Ross
Email: hayleyross@g.harvard.edu

Time: Tuesday 3-5pm
Location: TBA
Office hours: TBA or by appointment

1 Course Description

With models like ChatGPT on the rise and increasing hype about how they "learn language" just by consuming vast quantities of text, what is the role of linguistics? What do these models have to offer linguistics, and what do we have to offer them? This course aims to break down the barriers to using "computational" approaches, inspire students to consider computational and corpus-based approaches to problems, and critically consider the impact of large language models on linguistics and vice versa.

In this 6-week course, we'll start by exploring more traditional computational linguistics and show that computational approaches like probabilities, treebanks and word embeddings already have much to offer for principled explanations of interesting phenomena. Why do humans have trouble understanding sentences like *the horse raced past the barn fell*, even though there's only one syntactic tree for them? How did the meaning of *awful* – originally derived from *awe* – evolve to its current state and where is it headed? The second part of the course will focus on large language models like ChatGPT and discuss how we can apply linguistics to them and what the role of linguistics looks like with these models on the playing field. The class ends with students pursuing a project of their choice that can touch on any part of the course that resonated with them.

2 Prerequisites and Related Courses

This course has no prerequisites, though an introductory course in linguistics is highly recommended. In particular, no familiarity with programming, probability/statistics, or specifics of formal linguistic theories is required.

This course will have some overlap with COMPSCI 187 (Introduction to Computational Linguistics and Natural Language Processing), primarily in the area of context-free grammars and word embeddings, but will largely focus on different problems, and will cover only a small amount of the theoretical/mathematical background behind these approaches. Students interested in learning how to code the models that we discuss, apply more complex ("neural") models and/or apply these methods to practical tasks are encouraged to also take COMPSCI 187 (also offered in Fall 2023).

3 Schedule

Each class will be divided into two parts. The first part will present the new materials in a standard class format. The second part will be devoted to exercises in pairs or small groups. You will need to bring a laptop or tablet for most of the exercises.

Some weeks will have a required reading. Please read that *before* class that week. Conversely, you don't need to read the papers listed under 'Further Reading': they're just there in case you're interested in doing your final project or other research in that area.

Week 1:	*Introduction to probability and (probabilistic) context-free grammars*	
	Required reading:	none
	Further reading:	Goldsmith (2007) [for linguists]; Jurafsky (2003)
Week 2:	*Treebanks, counting, and garden path sentences*	
	Required reading:	Jurafsky & Martin, Section 17.1-17.3 and Section C.1
	Further reading:	Penn Treebank: Taylor, Marcus, and Santorini (2003) Dative: Bresnan et al. (2007), Rappaport Hovav and Levin (2008) Garden paths: Levy (2013)
		Assignment 1 released
Week 3:	*Distributional semantics*	
	Required reading:	Jurafsky & Martin, Section 6.1-6.4 (but skip the math in 6.4)
	Further reading:	Boleda (2020) [review article], Hamilton, Leskovec, and Jurafsky (2016), Marelli and Baroni (2015), Alammar (2019b), Mikolov et al. (2013)
		Assignment 1 due, Assignment 2 released
Week 4:	*Language models and surprisal*	
	Required reading:	Sections 1-2 and 5 of Linzen and Baroni (2021)
	Further reading:	Jurafsky & Martin, Chapter 10 (skim it), Part 1 of Alammar, 2019a, Belinkov and Glass (2019), Warstadt et al. (2020), Hu, Gauthier, et al. (2020), Hu, Levy, et al. (2023)
		Assignment 2 due, Assignment 3 released
Week 5:	*Linguistics vs. language models*	
	Required reading:	One of the following (to be divided up): Baroni (2022), Piantadosi (2023), Kodner, Payne, and Heinz (2023), Leivada and Grohmann (2023), Mahowald et al. (2023), McCloskey (1991)
		Assignment 3 due
Week 6:	*Bias & harm; project pitches*	
	Required reading:	none
	Further reading:	Jurafsky & Martin, Chapter 6.11; Papakyriakopoulos et al. (2020), TBD

4 Assignments and Grading

- **Participation (20%)** – attendance and involvement in in-class discussions and exercises
- **Assignments (45%)** – three short assignments, handed out in weeks 2, 3 and 4 and due the following week. Each assignment will build upon the in-class exercise, and counts for 15% of the overall grade. Some assignments will encourage you to come up with mini-research questions of your own; these will be graded based on the questions themselves and the steps taken to investigate them, regardless of the actual outcome. You have one late day which you may spend on any one assignment without any need for explanation.
- **Final project (35%)** – a short paper investigating the student's topic of choice, or an extension of one of the class topics, either using computational or probabilistic methods or critically engaging with them. (Students with further experience should feel free to use approaches not covered in class – please discuss this with me in advance.)

5 Textbook

We will be using the textbook 'Speech and Language Processing' by Jurafsky & Martin for part of the class. This textbook is available online; you do not need to purchase anything. All readings (including the relevant chapters of Jurafsky & Martin) will be available on Canvas.