Computational Seminar: Meaning Unmoored

Ling 7710, Spring 2021
Thurs 11:20-1:45, Zoom

Instructor: Marten van Schijndel
Email: mv443@cornell.edu
Office: Zoom
Office Hours: Mon & Wed 10-11
& By Appointment

Course Website: canvas.cornell.edu
Mailbox: 203 Morrill Hall
Office Phone: (607) 255-1105

Course Description: Neural networks have enabled amazing progress in natural language processing over the past decade. However, their linguistic representations have been repeatedly shown to simply exploit shallow statistical heuristics rather than actually learning the underlying linguistic patterns. This has sparked an active debate in natural language processing as to whether or not these models will ever be able to learn linguistic meaning, especially when most language models are trained solely on text data. In this seminar we will read and discuss a number of papers on cognitive theories of meaning, computational models of meaning, and statistical learning to look outside natural language processing for solutions to this dilemma.

Grading: This course revolves around discussions of research, so participation is a large component of the grade in this course. To facilitate discussion, students will be required to post responses to each week's reading to canvas the day before each class. Students are also expected to present relevant research papers or their own ongoing research. At the end of the course, students are expected to write a project proposal for a computational linguistics or cognitive science conference paper.[1]

Reading responses	10%								
Presentations	20%	105–97	A+	89.9–87	B+	79.9–77	C+	69.9–67	D+
Participation/discussion	20%	96.9–93	A	86.9–83	B	76.9–73	C	66.9–60	D
Final project proposal	50%	92.9–90	A-	82.9–80	B-	72.9–70	C-	Below 60	F

Reading Responses: Each person will need to submit reading responses via Canvas at least 24 hours in advance of the course. These need not be long, but they should demonstrate that you completed the reading and thought about it. They should contain a summary of the paper and a question to discuss in class. Five responses can be missed without penalty.

Presentations: During the course each student must present ongoing work which is relevant to the class, or they can instead present a paper relevant to the class.

Participation: Students are expected to engage with one another in detailed discussion of the research. Students should be informed and present during discussions and should regularly contribute. Astute, detailed contributions made as reading responses can stand in for constant speech during discussion periods, but students should each contribute to the course discussions in some way and should occasionally contribute verbally.

Final project proposal: Students will submit a project proposal or project write up at the end of the course. The project could explore the meaning representations of a pretrained language

[1] Alternatively, students can submit a draft of an actual conference paper developed during this course.

model (model probing), modify the meaning representations of a pretrained language model in accordance with the course papers (model augmentation), explore the data features that give rise to certain meaning representations in language models (data probing), or possibly use psycholinguistic experimentation to explore meaning representations in humans (human probing).

- Project topic proposed to me by March 25 (5%)
- In-person discussion with me about the project by April 8 (10%)
- Project proposal/paper (4-5 pages) to me by May 18
  - Research question (15%)
  - Background literature (20%)
  - Experimental design (20%)
  - Explanation of needed resources (e.g., sample size calculation) (10%)
  - Explanation of possible findings (what would each possible outcome mean?) (20%)

Final project proposals can be submitted individually or in pairs. Each student must contribute to the project design and the writing, and each group member must document the involvement of each student in the group in the comment section of the final project area of canvas.
Please write the proposal using the CogSci LaTeX template. Refer to Overleaf's 30 minute tutorial if you are new to LaTeX.

Special Accommodations: Please give me any Student Disability Services (SDS) accommodation letter as early as possible so that I can arrange for needed academic accommodations. If you need an immediate accommodation, please speak with me after class or email me and/or SDS at `sds_cu@cornell.edu`. SDS is located on level 5 of Cornell Health, 110 Ho Plaza, 607-254-4545, `https://sds.cornell.edu/`

Academic Integrity: Please don't cheat, and please cite people whose work you use. I am **required** to follow-up on suspected violations of Cornell's *Code of Academic Integrity*. I encourage students to watch this video to learn more about what constitutes cheating and why it matters: `https://cheatingvideo.provost.cornell.edu/`

If you have questions about issues of academic integrity, please see me.

Reading

This is an initial list of papers you might read to gain additional background for our discussions in the course. Students are encouraged to suggest additional papers that would be both relevant to this seminar and relevant to their research.

Optional Reading

- Bender & Koller (2020). Climbing towards NLU. ACL.
- Bisk et al (2020). Experience Grounds Language. EMNLP.
- Reitman and Bower (1973). Storage and later recognition of exemplars of concepts. Cognitive Psychology.
- Rudin (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence.
- Zador (2019). A Critique of Pure Learning. Nature.