CSCI-499: Natural Language for Interactive AI
Units: 4
Fall 2022 MoWe 10am-11:50am
Location: SOS B37

Instructor: Jesse Thomason
Office: SAL 244
Office Hours:
SAL 244
Tuesdays, 10am-11am
Contact Info: jessetho@usc.edu; Slack preferred

TA: Anthony Liang
Office Hours:
RTH 4th Floor Lounge
Tuesdays, 11am-12pm
Contact Info: aliang80@usc.edu; Slack preferred

Course Description
This course will explore how natural language can serve as an interaction medium between users and AI agents. To that end, the course will cover topics in natural language processing, computer vision, and machine learning. Students will become familiar with concepts and methods in natural language processing and linguistics such as syntactic and semantic parsing, structure prediction, distributional semantics and compositional semantics. Students will delve into the intersection of planning and search-oriented machine learning algorithms with such language understanding techniques and paradigms.

Learning Objectives

1. Gain familiarity with a diverse set of NLP concepts and techniques.
2. Gain familiarity with planning formalisms for learning-based AI agents.
3. Become familiar with and develop new techniques and applications for language as an interface between AI agents and human users.

Prerequisite(s): CSCI-270 && (CSCI-360 || CSCI-467 || equivalent experience)

Recommended Preparation: Fluency with python programming.

Course Notes

Lecture notes will be made available online after each class.

Required Readings and Supplementary Materials

All reading material will be posted on the course web page at the beginning of the course. All reading material will be freely and publicly available online.

Course Organization Fall 2022

Course communication planned via slack:

- Workspace: USC Viterbi School of Engineering Classes
  - uscviterbiclass.slack.com
  - Channel: #fall22-csci-499-30052

Course documents shared via Google Drive. Links:

- Course gDrive: https://drive.google.com/drive/folders/1Cn7nUqe1MMilSYTtK0C50DN7D-ab9HEQ?usp=sharing

- Coding Assignments Github: https://github.com/GLAMOR-USC/CSCI499_NaturalLanguageforInteractiveAI

Description and Assessment of Assignments

Paper reviews
The course will explore the course topics through a series of assigned readings in the form of research papers (and book chapters). Students will select one paper from a set of papers for each review cycle and submit a one page review of it as homework. There will be 8 (+1 optional) such paper reviews assigned through the semester. Reviews will be assessed based on answering the following 5 questions (based on clarity and correctness). In most cases, each question warrants at minimum a paragraph to answer, sometimes more. There are a few papers that are less technical for which these questions aren't as good of a fit; feel free to reach out to the instructor if you feel what you're writing about falls into that category.

1. What is the main problem/task addressed by the paper?
2. What was done before, and how does this paper improve on it?
3. What is the one cool technique/idea/finding that was learned from this paper?
4. What part of the paper was difficult to understand?
5. What generalization or extension of the paper could be done?

Coding Assignments
Students will explore course concepts in detail through coding assignments that require implementing core concepts in PyTorch and training and executing models against real data.

Class Presentation

Students will individually present a research paper to the class. The paper can be selected from a list of class papers from any review section, or outside of the selected list with instructor permission. Each presentation will be 10 mins total: 8 minutes for the presentation with 2 minutes for questions. The presentations will be graded on clarity, completeness, and presentation style.
What should you try to convey?

- The **main takeaway** you got from the paper; don't build suspense, drop the main information on us right away
- A brief **motivation** which can include some **background**; tell the audience why the paper needed to be written
- The **key insight** that overcomes the problems presented in the motivation with what was done before this paper
- The **technical details** behind the insight; go deep for a few slides (~40-50% of your speaking time)
- **Conclude** by summing it up: main takeaway driven by the key insight that resolved the problem presented in the motivation, with a reminder of technical detail/challenges

gDrive to upload slides:
https://drive.google.com/drive/folders/1SQoxsfhaPFYpwy48Q_vES16SvqVJk8Z6

gSheet to coordinate paper assignments:
https://docs.google.com/spreadsheets/d/1Uc9qvPSYbeUXF17j2pJYsU35qodleB2cdEuduPCkHrY/edit#gid=0

Semester Project
Students will work individually or in pairs to carry out a class project. The focus of the class project can be research-focused or application-focused. A research-focused project will develop models and analyze data of an existing problem in NLP, or formulate a new problem altogether. An

application-focused project will train (possibly only fine-tuning) and deploy NLP models to new application areas, while not necessarily developing any novel research question to be answered. Students will leverage tools, concepts, and techniques presented in the class. The project involves identifying a communication or exploration need that language could resolve, data sources available to inform the problem and method, and the techniques needed to approach it. The grade for the project will be based on the successful completion of the agreed upon project objectives. The deliverables include a project proposal (1-2 pages single space), a mid project report (4-8 pages single space), final presentation (15 minutes talk; 5 minutes for questions) and a final report (~8 pages single space for the main document, up to 15 with appendix/figures). Reports will be graded based on clarity, and completeness. The project is total 40% of final grade with the following internal breakdown:

Proposal (5%) [12.5 / 100 total pts]
Project Mid Report (10%) [25 / 100 total pts]
Final Presentation (10%) [25 / 100 total pts]
Project/Final Paper (15%) [37.5 / 100 total pts]

The project proposal (1-2 pages single space; Arial or LaTeX default font, no smaller than 11pt in any place) should outline the type of project (research-focused or application-focused), and then answer the following questions clearly in a sentence and/or a few paragraphs each, as appropriate.

- What are you trying to do? Articulate your objectives using absolutely no jargon.
- How is it done today, and what are the limits of current practice?

- What is new in your approach and why do you think it will be successful?
- Who cares? If you are successful, what difference will it make?
- What are the risks? What could go wrong, and how will you pivot early on if that happens?
- How much will it cost? That is, what resources will you need in terms of time and computation? Are these reasonable for a semester and what access you have?
- What are the midterm and final “exams” to check for success? That is, what are your milestones?

The project midterm report (about 5 pages single space; Arial or LaTeX default font, no smaller than 11pt in any place) should be broken into the following sections and cover requested content. If you are interested in writing your midterm and final report like a conference paper, you can build on the ACL Style Files and Templates.

- Introduction:
    - What are you trying to do? Articulate your objectives using absolutely no jargon.
    - Who cares? If you are successful, what difference will it make?
- Related Work:
    - How is it done today, and what are the limits of current practice?
- Method:
    - What is new in your approach and why do you think it will be successful?
    - How will you conduct experiments? What hypothesis will these experiments test? How will you evaluate the quality of the results?
- Preliminary Results:
    - Look back over the milestones you outlined in your proposal; did you hit your midterm exam to check for success?

    - What have you discovered so far? Did anything go the way you planned? Did anything go wrong in a way you have to change plans for? What are your next steps?
- Timeline to final report:
    - Provide a detailed timeline, including labor breakdown, of what your team will do between now and the final report.

The project final presentation should be 10-12m in length per group, leaving 3-5m for questions and speaker turnover. The presentations will be graded on clarity, completeness, and presentation style. Each member of the group should present for about equal time.
What should you try to convey?

- The **main takeaway** you have reached or aim to reach with your project; don't build suspense, drop the main information on us right away.
- A brief **motivation** which can include some **background**; tell the audience why the project needed to be done.
- The **key insight** that overcomes the problems presented in the motivation with what was done before this project.
- The **technical details** behind the insight; go deep for a few slides (~40-50% of your speaking time).
- **Conclude** by summing it up: main takeaway driven by the key insight that resolved the problem presented in the motivation, with a reminder of technical detail/challenges, and an overview of what remains to be done, if anything, before the final report, as well as what you hypothesize the outcomes of those experiments will be.

The project final report (no more than 8 pages single space; Arial or LaTeX default font, no smaller than 11pt in any place; up to 15 pages with appendix and additional figures) should be broken into the following sections and cover requested content. If you are interested in writing your

midterm and final report like a conference paper, you can build on the ACL Style Files and Templates.

- Introduction:
    - What were you trying to do? Articulate your objectives using absolutely no jargon.
    - Who cares? If you were successful, what difference would it make?
- Related Work:
    - How is it done today, and what are the limits of current practice?
- Method:
    - What is new in your approach and why did you think it would be successful?
    - How did you conduct experiments? What hypotheses did those experiments test? How did you evaluate the quality of the results?
- Results:
    - What are the results of your experiments? What did you discover? Were you able to support your hypotheses?
- Conclusions and Future Work:
    - What are the main takeaways of your project? What would you do next if you wanted to keep working in this space? What new questions can you formulate given the work this semester?

Grading Breakdown

Assignments	**Points**	**% of Grade**
Class participation	100	10%
Paper reviews (x8)	12.5 x 8	10%

Coding Assignments (x3) + Discussions/Participation	30 x 3 + 10	25%
Class Paper Presentation	100	15%
Final Project	100	40%

Final grades will be binned into letters (e.g., A+, A, etc.) with adjustments made for overall point distributions.

Assignment Submission Policy
Upload/email to be decided, encouraged by 6pm (aka “quittin’ time”) on the due date; accepted until 11:59pm on the due date. Blackboard “Assignments” will be created through which to turn stuff in.

Grading Timeline
Grades will be provided within 2 weeks of submission of the respective assignment.

Additional Policies
This is a discussion-based course, hence consistent attendance is expected. Missed classes with a valid excuse are allowed. Class participation will be scored based on engagements in course discussions.

The course will allow for a budget of 5 Late Day Tokens per student. These tokens can be expended on project deliverables (NOT presentations) to extend the deadline for a student. These tokens should be used with no justification or explanation for taking the late time required (i.e., you do not need to explain your reasoning to me). Going over budget (e.g., turning things in late with no Late Day Tokens to expend) will incur grade penalties of 5% per day late. For project teams, each member must expend a Late Day Token for a single day extension to a project-related deadline (i.e., for a team of 2 to push back their midterm report deadline, both team

members must expend a Late Day Token). *There are no refunds for late days*: unused late days cannot be converted into credit of any kind. Use them!

Project Schedule:
Week 5: Proposals due Thu Sep 22
Week 10: Midterm Reports Thu Oct 27
Week 14-15: Project Presentations given in class
Final Report: Due Monday, December 12 [USC-scheduled Final Exam Day]

Course Schedule:
Everything in the schedule is subject to change *except* the Project Schedule. Topics/Daily Activities especially are a very vague outline that will become concrete as we move through the weeks and identify what needs to be covered. Homework and review deadlines ("Deliverables") are more solidly set. The Project Schedule dates are *unlikely* to change.

	Topics/Daily Activities	**Deliverables ›**
Week 1 Mon Aug 22	Lecture 0: Administrivia Lecture 1: Formalizing Dialogue	
Wed Aug 24	Lecture 1: Formalizing Dialogue Lecture 2: Text Classification	Coding assignment 1 *released* to class
Week 2 Mon Aug 29	Lecture 2: Text Classification	

Wed Aug 31	Lecture 2: Text Classification	Paper review 1 DUE Thu Sep 1
Week 3 Mon Sep 5	NO CLASS [Labor Day]	NO CLASS
Wed Sep 7	Lecture 2: Text Classification Lecture 3: Language Modeling	Coding assignment 1: text classification DUE Thu Sep 8
Week 4 Mon Sep 12	Lecture 3: Language Modeling	
Wed Sep 14	Coding Assignment 1 Debrief and Discussion	Paper Review 2 DUE Thu Sep 15
Week 5 Mon Sep 19	Lecture 3: Language Modeling	
Wed Sep 21	Lecture 4: Word Embeddings	Paper Review 3 DUE Thu Sep 22 Project Proposal DUE Thu Sep 22
Week 6 Mon Sep 26	Lecture 4: Word Embeddings	
Wed Sep 28	Lecture 4: Word Embeddings Detour Lecture: Language Grounding	Paper Review 4 DUE Thu Sep 29 Coding assignment 2 *released* to class

Week 7 Mon Oct 3	Student Paper Presentations	Paper Presentations DUE
Wed Oct 5	Student Paper Presentations	Optional Paper Review DUE Thu Oct 6 Paper Presentations DUE
Week 8 Mon Oct 10	NO CLASS [Indigenous People's Day]	NO CLASS
Wed Oct 12	Detour Lecture: Language Grounding Models + Coding Assignment 2 open discussion ahead of turn in	Paper Review 5 DUE Thu Oct 13 Coding assignment 2: Word Embeddings DUE Thu Oct 13
Week 9 Mon Oct 17	Detour Lecture: Language Grounding Lecture 5: Sequence-to-Sequence Models	
Wed Oct 19	Lecture 5: Sequence-to-Sequence Models	Paper Review 6 DUE Thu Oct 20
Week 10 Mon Oct 24	Lecture 5: Sequence-to-Sequence Models	

Wed Oct 26	Lecture 5: Sequence-to-Sequence Models Coding Assignment 2 Debrief and Discussion	Project mid report DUE Thu Oct 27 Homework 3 RELEASED Fri Oct 28
Week 11 Mon Oct 31	Coding Assignment 2 Debrief and Discussion Lecture 6: NLP for Interaction with a World	
Wed Nov 2	Lecture 6: NLP for Interaction with a World	Paper Review 7 DUE Thu Nov 3
Week 12 Mon Nov 7	Lecture 6: NLP for Interaction with a World	
Wed Nov 9	Lecture 6: NLP for Interaction with a World	Paper Review 8 DUE Thu Nov 10
Week 13 Mon Nov 14	Lecture 6: NLP for Interaction with a World Models + Coding Assignment 3 open discussion ahead of turn in	Coding assignment 3: Seq2Seq DUE Tue Nov 15
Wed Nov 16	Lecture 7: NLP for Embodied Interaction	
Week 14 Mon Nov 21	GLAMOR Lab presentations on NLP for Interactive AI	

	-Sign Language Recognition: Lee Kezar -Continual Learning: Tejas Srinivasan -LLM Prompting: Ying-Yun Chang -Prompting for Control: Ishika Singh	
Wed Nov 23	NO CLASS [Day before Thanksgiving holiday]	NO CLASS
Week 15 Mon Nov 28	**Project Presentations** ***In person session*** [NOTE: last day of in-person class for us]	**Project Presentations**
Wed Nov 30	**Project Presentations** ***Virtual session***	**Project Presentations**
FINAL	**Project Report**	**Project Report** DUE Monday, December 12 [USC-scheduled final]

Paper Review Candidates:

Below are suggested papers for each review assignment. If you choose to review a paper outside this area, you must first obtain instructor approval for the candidate outside paper. Papers may be added to these sections through the semester, but we will avoid removing them.

Paper Review 1

- Logic and Conversation. Paul Grice. Syntax and Semantics 3: Speech Acts. Academic Press 1975.

- Executing Instructions in Situated Collaborative Interactions. Alane Suhr, Claudia Yan, Jacob Schluger, Stanley Yu, Hadi Khader, Marwa Mouallem, Iris Zhang, Yoav Artzi. EMNLP 2019.
- Vision-and-Dialog Navigation. Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer. CoRL 2019.
- Collaborative dialogue in Minecraft. Anjali Narayan-Chen, Prashant Jayannavar, Julia Hockenmaier. ACL 2019.
- TEACh: Task-driven Embodied Agents that Chat. Aishwarya Padmakumar, Jesse Thomason, Ayush Shrivastava, Patrick Lange, Anjali Narayan-Chen, Spandana Gella, Robinson Piramithu, Gokhan Tur, and Dilek Hakkani-Tur. AAAI 2022.
- Recipes for building an open-domain chatbot. Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric M. Smith, Y-Lan Boureau, Jason Weston. ArXiv, 2020.

Paper Review 2

- Efficient Estimation of Word Representations in Vector Space. Tomas Mikolov, Kai Chen, Greg Corrado, Jeffrey Dean. ICLR 2013.
- Deep contextualized word representations. Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, Luke Zettlemoyer. NAACL 2018.
- Attention Is All You Need. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, Illia Polosukhin. NeurIPS 2017.
- A Structured Vector Space Model for Word Meaning in Context. Katrin Erk and Sebastian Padó. EMNLP 2008.
- ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks. Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. CVPR 2020.
- Grounding Language in Play. Corey Lynch and Pierre Sermanet. 2020.
- Learning Language-Conditioned Robot Behavior from Offline Data and Crowd-Sourced Annotation. Suraj Nair and Eric Mitchell and Kevin Chen and Brian Ichter and Silvio Savarese and Chelsea Finn. CoRL 2021.

Paper Review 3

- The Curious Case of Neural Text Degeneration. Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, Yejin Choi. International Conference on Learning Representations (ICLR), 2020.
- Synthetic and Natural Noise Both Break Neural Machine Translation. Yonatan Belinkov, Yonatan Bisk. International Conference on Learning Representations (ICLR), 2018.
- Distributed Representations of Words and Phrases and their Compositionality. Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, Jeffrey Dean. Neural Information Processing Systems (NeurIPS) 2013.
- Learning to Interpret Natural Language Navigation Instructions from Observations. David L. Chen and Raymond J. Mooney. AAAI 2011.
- Vision-and-Language Navigation: Interpreting visually-grounded navigation instructions in real environments. Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, Anton van den Hengel. CVPR 2018.

- Improving Vision-and-Language Navigation with Image-Text Pairs from the Web. Arjun Majumdar and Ayush Shrivastava and Stefan Lee and Peter Anderson and Devi Parikh and Dhruv Batra. ECCV 2020.
- Shifting the Baseline: Single Modality Performance on Visual Navigation & QA. Jesse Thomason, Daniel Gordon, and Yonatan Bisk. North American Chapter of the Association for Computational Linguistics (NAACL), 2019.

Paper Review 4

- GloVe: Global Vectors for Word Representation. Jeffrey Pennington, Richard Socher, Christopher Manning. EMNLP 2014.
- Sense Embedding Learning for Word Sense Induction. Linfeng Song, Zhiguo Wang, Haitao Mi, Daniel Gildea. *SEM 2016.
- Enriching Word Vectors with Subword Information. Piotr Bojanowski, Edouard Grave, Armand Joulin, Tomas Mikolov. TACL 2017.
- VQA: Visual Question Answering. Stanislaw Antol and Aishwarya Agrawal and Jiasen Lu and Margaret Mitchell and Dhruv Batra and C. Lawrence Zitnick and Devi Parikh. ICCV 2015.
- Information Maximizing Visual Question Generation. Ranjay Krishna and Michael Bernstein and Li Fei-Fei. CVPR 2019.
- Grounded Situation Recognition. Sarah Pratt and Mark Yatskar and Luca Weihs and Ali Farhadi and Aniruddha Kembhavi. ECCV 2020.

Paper Review 5

- Entailment Semantics Can Be Extracted from an Ideal Language Model. William Merrill, Alex Warstadt, Tal Linzen. CoNLL 2022.
- Multi-Modal Word Synset Induction. Jesse Thomason and Raymond J. Mooney. IJCAI 2017.
- Does Pretraining for Summarization Require Knowledge Transfer? Kundan Krishna, Jeffrey Bigham, Zachary C. Lipton. EMNLP 2021.
- Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering. Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, Lei Zhang. CVPR 2018.
- Learning Transferable Visual Models From Natural Language Supervision. Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, Ilya Sutskever. 2021.
- ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. Jiasen Lu, Dhruv Batra, Devi Parikh, Stefan Lee. NeurIPS 2019.
- ViLT: Vision-and-Language Transformer Without Convolution or Region Supervision. Kim, Wonjae and Son, Bokyung and Kim, Ildoo. ICML 2021.
- MERLOT: Multimodal Neural Script Knowledge Models. Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, Yejin Choi. NeurIPS 2021.

Paper Review 6

- Is Reinforcement Learning (Not) for Natural Language Processing?: Benchmarks, Baselines, and Building Blocks for Natural Language Policy Optimization. Rajkumar Ramamurthy, Prithviraj Ammanabrolu, Kianté Brantley, Jack Hessel, Rafet Sifa, Christian Bauckhage, Hannaneh Hajishirzi, Yejin Choi. arXiv, 2022.
- Downstream Datasets Make Surprisingly Good Pretraining Corpora. Kundan Krishna, Saurabh Garg, Jeffrey P. Bigham, Zachary C. Lipton. Arxiv 2022.
- VIMA: General Robot Manipulation with Multimodal Prompts. Yunfan Jiang, Agrim Gupta, Zichen Zhang, Guanzhi Wang, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu, Linxi Fan. arXiv, 2022.
- https://larel-ws.github.io/accepted-papers/ [any paper you like from these proceedings]
- PixL2R: Guiding Reinforcement Learning Using Natural Language by Mapping Pixels to Rewards. Prasoon Goyal, Scott Niekum, Raymond J. Mooney. 2020.
- Inferring Rewards from Language in Context. Jessy Lin, Daniel Fried, Dan Klein, Anca Dragan. ACL 2022.

Paper Review 7

- A New Path: Scaling Vision-and-Language Navigation with Synthetic Instructions and Imitation Learning. Aishwarya Kamath, Peter Anderson, Su Wang, Jing Yu Koh, Alexander Ku, Austin Waters, Yinfei Yang, Jason Baldridge, Zarana Parekh. arXiv, 2022.
- Mapping Navigation Instructions to Continuous Control Actions with Position Visitation Prediction. Valts Blukis and Dipendra Misra and Ross A. Knepper and Yoav Artzi. CoRL 2018.
- Sim-to-Real Transfer for Vision-and-Language Navigation. Peter Anderson and Ayush Shrivastava and Joanne Truong and Arjun Majumdar and Devi Parikh and Dhruv Batra and Stefan Lee. CoRL 2020.
- Asking for Help Using Inverse Semantics. Stefanie Tellex and Ross Knepper and Adrian Li and Daniela Rus and Nicholas Roy. RSS 2014.
- CLIPort: What and Where Pathways for Robotic Manipulation. Mohit Shridhar, Lucas Manuelli, Dieter Fox. CoRL 2021.
- https://say-can.github.io/ [Post evolves faster than the paper, in this case.]

Paper Review 8

- Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data. Emily M. Bender and Alexander Koller. ACL 2020.
- Experience Grounds Language. Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. EMNLP 2020.
- ProgPrompt: Generating Situated Robot Task Plans using Large Language Models. Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, Jesse Thomason, and Animesh Garg. arXiv, 2022.
- CLiMB: A Continual Learning Benchmark for Vision-and-Language Tasks. Tejas Srinivasan, Ting-Yun Chang, Leticia Leonor Pinto Alva, Georgios Chochlakis, Mohammad Rostami, and Jesse Thomason. Neural Information Processing Systems (NeurIPS), 2022.

- The World of an Octopus: How Reporting Bias Influences a Language Model's Perception of Color. Cory Paik, Stéphane Aroca-Ouellette, Alessandro Roncone, Katharina Kann. EMNLP 2021.
- Unnatural Language Processing: Bridging the Gap Between Synthetic and Natural Language Data. Alana Marzoev, Samuel Madden, M. Frans Kaashoek, Michael Cafarella, Jacob Andreas. 2020.
- Towards Ecologically Valid Research on Language User Interfaces. Harm de Vries and Dzmitry Bahdanau and Christopher Manning. 2020.
- Dynabench: Rethinking Benchmarking in NLP. Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, Adina Williams. NAACL 2021.

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 Part B, Section 11, "Behavior Violating University Standards" policy.usc.edu/scampus-part-b. Other forms of academic dishonesty are equally unacceptable. See additional information in SCampus and university policies on scientific misconduct, policy.usc.edu/scientific-misconduct.

Support Systems:

Counseling and Mental Health - (213) 740-9355 – 24/7 on call
studenthealth.usc.edu/counseling
Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

National Suicide Prevention Lifeline - 1 (800) 273-8255 – 24/7 on call
suicidepreventionlifeline.org
Free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week.

Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-9355(WELL), press "0" after hours – 24/7 on call
studenthealth.usc.edu/sexual-assault
Free and confidential therapy services, workshops, and training for situations related to gender-based harm.

Office of Equity and Diversity (OED) - (213) 740-5086 | Title IX – (213) 821-8298

equity.usc.edu, titleix.usc.edu
Information about how to get help or help someone affected by harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants.

Reporting Incidents of Bias or Harassment - (213) 740-5086 or (213) 821-8298
usc-advocate.symplicity.com/care_report
Avenue to report incidents of bias, hate crimes, and microaggressions to the Office of Equity and Diversity |Title IX for appropriate investigation, supportive measures, and response.

The Office of Disability Services and Programs - (213) 740-0776
dsp.usc.edu
Support and accommodations for students with disabilities. Services include assistance in providing readers/notetakers/interpreters, special accommodations for test taking needs, assistance with architectural barriers, assistive technology, and support for individual needs.

USC Campus Support and Intervention - (213) 821-4710
campussupport.usc.edu
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

Diversity at USC - (213) 740-2101
diversity.usc.edu

Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

USC Emergency - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call
dps.usc.edu, emergency.usc.edu
Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

USC Department of Public Safety - UPC: (213) 740-6000, HSC: (323) 442-120 – 24/7 on call
dps.usc.edu
Non-emergency assistance or information.