

# Psych/CSCI 626: Text as Data

Morteza Dehghani

Fall 2024

E-mail: [mdehghan@usc.edu](mailto:mdehghan@usc.edu)  
Office Hours: Mon 10-12  
Office: SGM 607

Web: [mola-lab.org](http://mola-lab.org)  
Class Hours: Th 2-6pm  
Class Room: SGM 911

---

## Course Description

Text as Data focuses on applications of natural language processing (NLP), guided by psychological theories, for identifying various social and cognitive properties evident in human-related data. In this course, we will survey the application of various NLP techniques, with an emphasis on Large Language Models, for investigating various aspects of human cognition. The intended audience for this course is psychology and computer science PhD students, and more broadly graduate students in who are interested in using NLP for data analysis.

## Learning Objectives

This course is designed to survey the current state of research in automated language analysis for making inferences about various aspects of human cognition. It's important to note that this class does not aim to teach NLP or psychology, but rather to explore how established NLP methods can be used to study various social and cognitive phenomena. Optional reading material will be provided for students who are unfamiliar with the topics discussed.

- **Prerequisite(s):** Instructor permission
- **Recommended Preparation:** For Non-Engineering majors: Psych 625 or a similar course, for Engineering students: CSCI 544 or a similar course.

## Course Notes

Students are not allowed to use laptops or smartphones during class, unless used for class presentation. Assignments will be posted on Bright Horizons.

## Required Books

- Pennebaker, J. (2011). *The secret life of pronouns: What our words say about us*. New York, NY: Bloomsbury

## Description and Assessment of Assignments

1. Reaction paragraphs. Students are asked to write a short note, about two paragraphs in length, about their reaction to the reading assignments of the week. These can be comments about the subject area, or a critique of a particular theory or experiment. I will read these paragraphs carefully before each class, and will use them to guide the discussion in class. Simply reading the first page of a paper and writing a summary of it will not count as a reaction paragraph. Using an LLM to write your reaction paragraphs will serve no educational purpose.
2. Class Project. This class is project oriented, and group-based. The goal of the project is for students to get experience working in interdisciplinary groups to tackle specific social scientific problems, and bring together theory from the social sciences and NLP techniques from computer science to tackle that problem. This will include a project proposal presentation, three project update presentations, final project presentation, and a report. For project proposals, students will present a problem and a data collection method and/or dataset for which they want to analyze. Each presentation should be about 10-15mins. The goal of the project update presentations is to inform the class about the state of the project and brainstorm with other students on how to solve the remaining issues. Each update presentation should be around 10 minutes. For the final project presentation, each student/group will give a 15-20min presentation on their project. Students are expected to spend at least 80 hours working on their final project. The project report will be around 20 pages.

## Grading Policy

- 30% Discussion Participation
- 30% Reaction Paragraphs
- 15% Project Status Updates
- 25% Final Project Presentation and Write up

## Assignment Submission Policy

All assignments are due on Thursdays at 9am. Assignments turned in any later than 9:10am will be considered late. Students will be allowed a total of four late days that can be used on the assignments. In exceptional circumstances, arrangements must be made in advance of the due date to obtain an extension. Once you have used up your four late days, one additional day late will result in a 25% reduction in the total score, two additional days late will yield a 50% reduction, and no credit will be given for three or more additional days late. Late days are in units of days, not hours, so using up part of a day uses up the whole day. The final project report, plus code used, will be due on the day of the final exam.

## Use of LLMs

You are welcome to utilize LLMs in this class to enhance the fluency of your writing. Additionally, LLMs can be helpful tools for debugging and refining your code. Learning to use AI is an emerging skill, but please keep in mind the following:

- AI tools are permitted to help you brainstorm topics or revise work you have already written.
- If you provide minimum-effort prompts, you will get low-quality results. You will need to refine your prompts to get good outcomes. This will take work.
- Proceed with caution when using AI tools and do not assume the information provided is accurate or trustworthy. If it gives you a number or fact, assume it is incorrect unless you either know the correct answer or can verify its accuracy with another source. You will be responsible for any errors or omissions provided by the tool. It works best for topics you understand.
- AI is a tool, but one that you need to acknowledge using. Please include a paragraph at the end of any assignment that uses AI explaining how (and why) you used AI and indicate/specify the prompts you used to obtain the results what prompts you used to get the results. Failure to do so is a violation of academic integrity policies.
- Be thoughtful about when AI is useful. Consider its appropriateness for each assignment or circumstance. The use of AI tools requires attribution. You are expected to clearly attribute any material generated by the tool used.

## Schedule and weekly learning goals

The schedule is tentative and subject to change.

### Week 01, 08/29: Introduction to Computational Social Sciences

- Lazer, D., Hargittai, E., Freelon, D., Gonzalez-Bailon, S., Munger, K., Ognyanova, K., and Radford, J. (2021). Meaningful measures of human society in the twenty-first century. *Nature*, 595(7866):189–196
- Wallach, H. (2018). Computational social science  $\neq$  computer science + social data. *Communications of the ACM*, 61(3):42–44
- Kennedy, B., Ashokkumar, A., Boyd, R. L., and Dehghani, M. (2022). Text analysis for psychology: Methods, principles, and practices. In Dehghani, M. and Boyd, R. L., editors, *Handbook of Language Analysis in Psychology*. Guilford Press, New York, NY

### Week 02, 09/05: Jamie

- Pennebaker, J. (2011). *The secret life of pronouns: What our words say about us*. New York, NY: Bloomsbury

**Week 03, 09/12:** Morteza's Favorites

- Park, G., Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Kosinski, M., Stillwell, D. J., Ungar, L. H., and Seligman, M. E. (2015). Automatic personality assessment through social media language. *Journal of personality and social psychology*, 108(6):934
- Voigt, R., Camp, N. P., Prabhakaran, V., Hamilton, W. L., Hetey, R. C., Griffiths, C. M., Jurgens, D., Jurafsky, D., and Eberhardt, J. L. (2017). Language from police body camera footage shows racial disparities in officer respect. *Proceedings of the National Academy of Sciences*, 114(25):6521–6526
- Rho, E. H., Harrington, M., Zhong, Y., Pryzant, R., Camp, N. P., Jurafsky, D., and Eberhardt, J. L. (2023). Escalated police stops of black men are linguistically and psychologically distinct in their earliest moments. *Proceedings of the National Academy of Sciences*, 120(23):e2216162120
- Jackson, J. C., Watts, J., Henry, T. R., List, J.-M., Forkel, R., Mucha, P. J., Greenhill, S. J., Gray, R. D., and Lindquist, K. A. (2019). Emotion semantics show both cultural variation and universal structure. *Science*, 366(6472):1517–1522
- Mooijman, M., Hoover, J., Lin, Y., Ji, H., and Dehghani, M. (2018). Moralization in social networks and the emergence of violence during protests. *Nature human behaviour*, 2(6):389

**Week 04, 09/19:** Debates

- Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., and Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28):7313–7318
  - Burton, J. W., Cruz, N., and Hahn, U. (2021). Reconsidering evidence of moral contagion in online social networks. *Nature Human Behaviour*, pages 1–7
  - Wang, S.-Y. N. and Inbar, Y. (2022). Re-examining the spread of moralized rhetoric from political elites: Effects of valence and ideology. *Journal of Experimental Psychology: General*
  - Candia, C., Atari, M., Kteily, N., and Uzzi, B. (2022). Overuse of moral language dampens content engagement on social media. *SocArXiv*
- Youyou, W., Yang, Y., and Uzzi, B. (2023a). A discipline-wide investigation of the replicability of psychology papers over the past two decades. *Proceedings of the National Academy of Sciences*, 120(6):e2208863120
  - Mottelson, A. and Kontogiorgos, D. (2023). Replicating replicability modeling of psychology papers. *Proceedings of the National Academy of Sciences*, 120(33):e2309496120

- Crockett, M., Bai, X., Kapoor, S., Messeri, L., and Narayanan, A. (2023). The limitations of machine learning models for predicting scientific replicability. *Proceedings of the National Academy of Sciences*, 120(33):e2307596120
- Youyou, W., Yang, Y., and Uzzi, B. (2023b). Reply to crockett et al. and mottelson and kontogiorgos: Machine learning’s scientific significance and future impact on replicability research. *Proceedings of the National Academy of Sciences*, 120(33):e2308195120

### Week 05, 09/26: Large Language Models I: Potential for Social Sciences

- Ziems, C., Held, W., Shaikh, O., Chen, J., Zhang, Z., and Yang, D. (2024). Can Large Language Models Transform Computational Social Science? *Computational Linguistics*, 50(1):237–291
- Rathje, S., Mirea, D.-M., Sucholutsky, I., Marjeh, R., Robertson, C. E., and Van Bavel, J. J. (2024). Gpt is an effective tool for multilingual psychological text analysis. *Proceedings of the National Academy of Sciences*, 121(34):e2308950121
- Ashokkumar, A., Hewitt, L., Ghezze, G., and Willer, R. (2024). Predicting results of social science experiments using large language models. *Under Review*
- Wang, A., Morgenstern, J., and Dickerson, J. P. (2024). Large language models cannot replace human participants because they cannot portray identity groups
- Abdurahman, S., Atari, M., Karimi-Malekabadi, F., Xue, M. J., Trager, J., Park, P. S., Golazizian, P., Omrani, A., and Dehghani, M. (2024). Perils and opportunities in using large language models in psychological research. *PNAS Nexus*, 3(7):pgae245
- Messeri, L. and Crockett, M. (2024). Artificial intelligence and illusions of understanding in scientific research. *Nature*, 627(8002):49–58

### Week 06, 10/03: Large Language Models II: Cognitive abilities

- Piantadosi, S. (2023). Modern language models refute chomsky’s approach to language. *Lingbuzz Preprint*, *lingbuzz*, 7180
- Webb, T., Holyoak, K. J., and Lu, H. (2023). Emergent analogical reasoning in large language models. *Nature Human Behaviour*, pages 1–16
  - Hodel, D. and West, J. (2024). Response: Emergent analogical reasoning in large language models
  - Mitchell, M. (2023). On analogy-making in large language models
  - Webb, J. (2024). Response to “on analogy-making in large language models”

- Binz, M. and Schulz, E. (2023). Using cognitive psychology to understand gpt-3. *Proceedings of the National Academy of Sciences*, 120(6):e2218523120
- Mitchell, M. and Krakauer, D. C. (2023). The debate over understanding in AI's large language models. *Proceedings of the National Academy of Sciences*, 120(13):e2215907120

**Week 07, 10/10:** Fall Recess

**Week 08, 10/17:** Large Language Models III & Project Proposal Presentations

- Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., and Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22
- Hua, W., Fan, L., Li, L., Mei, K., Ji, J., Ge, Y., Hemphill, L., and Zhang, Y. (2024). War and peace (waragent): Large language model-based multi-agent simulation of world wars
- Chuang, Y.-S., Goyal, A., Harlalka, N., Suresh, S., Hawkins, R., Yang, S., Shah, D., Hu, J., and Rogers, T. T. (2023). Simulating opinion dynamics with networks of llm-based agents. *arXiv preprint arXiv:2311.09618*
- Shanahan, M., McDonnell, K., and Reynolds, L. (2023). Role play with large language models. *Nature*, 623(7987):493–498
- Shen, S., Logeswaran, L., Lee, M., Lee, H., Poria, S., and Mihalcea, R. (2024). Understanding the capabilities and limitations of large language models for cultural commonsense. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics*

**Week 09, 10/24:** Large Language Models IV: Subjectivity

- Davani, A. M., Atari, M., Kennedy, B., and Dehghani, M. (2023). Hate speech classifiers learn normative social stereotypes. *Transactions of the Association for Computational Linguistics*, 11:300–319
- Davani, A., Díaz, M., Baker, D., and Prabhakaran, V. (2024). Disentangling perceptions of offensiveness: Cultural and moral correlates. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 2007–2021
- Orlikowski, M., Röttger, P., Cimiano, P., and Hovy, D. (2023). The ecological fallacy in annotation: Modelling human label variation goes beyond sociodemographics. *arXiv preprint arXiv:2306.11559*
- Sap, M., Swayamdipta, S., Vianna, L., Zhou, X., Choi, Y., and Smith, N. A. (2022). Annotators with attitudes: How annotator beliefs and identities bias toxic language detection. In *NAACL*

- Golazizian, P., Omrani, A., Ziabari, A. S., and Dehghani, M. (2024). Cost-efficient subjective task annotation and modeling through few-shot annotator adaptation. *arXiv preprint arXiv:2402.14101*

### Week 10, 10/31: Cognitive applications I & Project Update I

- Mitchell, T. M., Shinkareva, S. V., Carlson, A., Chang, K.-M., Malave, V. L., Mason, R. A., and Just, M. A. (2008). Predicting human brain activity associated with the meanings of nouns. *science*, 320(5880):1191–1195
- Dehghani, M., Boghrati, R., Man, K., Hoover, J., Gimbel, S. I., Vaswani, A., Zevin, J. D., Immordino-Yang, M. H., Gordon, A. S., Damasio, A., et al. (2017). Decoding the neural representation of story meanings across languages. *Human brain mapping*, 38(12):6096–6106
- Kumar, S., Sumers, T. R., Yamakoshi, T., Goldstein, A., Hasson, U., Norman, K. A., Griffiths, T. L., Hawkins, R. D., and Nastase, S. A. (2024). Shared functional specialization in transformer-based language models and the human brain. *Nature communications*, 15(1):5523
- Jamali, M., Grannan, B., Cai, J., Khanna, A. R., Muñoz, W., Caprara, I., Paulk, A. C., Cash, S. S., Fedorenko, E., and Williams, Z. M. (2024). Semantic encoding during language comprehension at single-cell resolution. *Nature*, pages 1–7

### Week 11, 11/07: Cognitive applications II

- Goldstein, A., Ham, E., Nastase, S. A., Zada, Z., Grinstein-Dabus, A., Aubrey, B., Schain, M., Gazula, H., Feder, A., Doyle, W., Devore, S., Dugan, P., Friedman, D., Brenner, M., Hassidim, A., Devinsky, O., Flinker, A., Levy, O., and Hasson, U. (2023). Correspondence between the layered structure of deep language models and temporal structure of natural language processing in the human brain. *bioRxiv*
- Card, N. S., Wairagkar, M., Iacobacci, C., Hou, X., Singer-Clark, T., Willett, F. R., Kunz, E. M., Fan, C., Vahdati Nia, M., Deo, D. R., et al. (2024). An accurate and rapidly calibrating speech neuroprosthesis. *New England Journal of Medicine*, 391(7):609–618
- Goldstein, A., Grinstein-Dabush, A., Schain, M., Wang, H., Hong, Z., Aubrey, B., Schain, M., Nastase, S. A., Zada, Z., Ham, E., et al. (2024). Alignment of brain embeddings and artificial contextual embeddings in natural language points to common geometric patterns. *Nature communications*, 15(1):2768
- Zada, Z., Goldstein, A., Michelmann, S., Simony, E., Price, A., Hasenfratz, L., Barham, E., Zadbood, A., Doyle, W., Friedman, D., et al. (2024). A shared model-based linguistic space for transmitting our thoughts from brain to brain in natural conversations. *Neuron*

**Week 12, 11/14:** Bias in NLP

- Garg, N., Schiebinger, L., Jurafsky, D., and Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644
- Charlesworth, T. E., Sanjeev, N., Hatzenbuehler, M. L., and Banaji, M. R. (2023). Identifying and predicting stereotype change in large language corpora: 72 groups, 115 years (1900–2015), and four text sources. *Journal of Personality and Social Psychology*
- Hovy, D. and Prabhumoye, S. (2021). Five sources of bias in natural language processing. *Language and Linguistics Compass*, 15(8):e12432
- Cheng, M., Durmus, E., and Jurafsky, D. (2023). Marked personas: Using natural language prompts to measure stereotypes in language models. *arXiv preprint arXiv:2305.18189*
- Omrani, A., Ziabari, A. S., Yu, C., Golazizian, P., Kennedy, B., Atari, M., Ji, H., and Dehghani, M. (2023). Social-group-agnostic bias mitigation via the stereotype content model. In *Proc. The 61st Annual Meeting of the Association for Computational Linguistics (ACL2023)*

**Week 13, 11/21:** Ethics & Project Update II

- Weidinger, L., Mellor, J., Rauh, M., Griffin, C., Uesato, J., Huang, P.-S., Cheng, M., Glaese, M., Balle, B., Kasirzadeh, A., et al. (2021). Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359*
- Andrews, M., Smart, A., and Birhane, A. (2024). The reanimation of pseudoscience in machine learning and its ethical repercussions. *Patterns*
- Skorburg, J. A. and Friesen, P. (2022). Ethical issues in text mining for mental health. In Dehghani, M. and Boyd, R. L., editors, *Handbook of Language Analysis in Psychology*. Guilford Press, New York, NY
- Bender, E. M. (2024). Resisting dehumanization in the age of “AI”. *Current Directions in Psychological Science*, 33(2):114–120
- Bender, E. M., Gebru, T., McMillan-Major, A., and Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, pages 610–623, New York, NY, USA. Association for Computing Machinery
- Wa Thiong’o, N. (1998). Decolonising the mind. *Diogenes*, 46(184):101–104
- Watch in class: Friends You Haven’t Met Yet

**Week 14, 11/28:** Thanksgiving Holiday



**Week 15, 12/05:** Final project presentations

---

## Statement on Academic Conduct and Support Systems

### Academic Integrity

The University of Southern California is a learning community committed to developing successful scholars and researchers dedicated to the pursuit of knowledge and the dissemination of ideas. Academic misconduct, which includes any act of dishonesty in the production or submission of academic work, compromises the integrity of the person who commits the act and can impugn the perceived integrity of the entire university community. It stands in opposition to the university's mission to research, educate, and contribute productively to our community and the world.

All students are expected to submit assignments that represent their own original work, and that have been prepared specifically for the course or section for which they have been submitted. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s).

Other violations of academic integrity include, but are not limited to, cheating, plagiarism, fabrication (e.g., falsifying data), collusion, knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

The impact of academic dishonesty is far-reaching and is considered a serious offense against the university. All incidences of academic misconduct will be reported to the Office of Academic Integrity and could result in outcomes such as failure on the assignment, failure in the course, suspension, or even expulsion from the university.

For more information about academic integrity see the student handbook or the Office of Academic Integrity's website, and university policies on Research and Scholarship Misconduct.

Please ask your instructor if you are unsure what constitutes unauthorized assistance on an exam or assignment, or what information requires citation and/or attribution.

### Students and Disability Accommodations

USC welcomes students with disabilities into all of the University's educational programs. The Office of Student Accessibility Services (OSAS) is responsible for the determination of appropriate accommodations for students who encounter disability-related barriers. Once a student has completed the OSAS process (registration, initial appointment, and submitted documentation) and accommodations are determined to be reasonable and appropriate, a Letter of Accommodation (LOA) will be available to generate for each course. The LOA must be given to each course instructor by the student and followed up with a discussion. This should be done as early in the semester as possible as accommodations are not retroactive. More information can be found at [osas.usc.edu](http://osas.usc.edu). You may contact OSAS at (213) 740-0776 or via email at [osasfrontdesk@usc.edu](mailto:osasfrontdesk@usc.edu).

### Support Systems

*Counseling and Mental Health - (213) 740-9355 - 24/7 on call*

Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

*988 Suicide and Crisis Lifeline - 988 for both calls and text messages - 24/7 on call*

The 988 Suicide and Crisis Lifeline (formerly known as the National Suicide Prevention Lifeline)

provides free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week, across the United States. The Lifeline is comprised of a national network of over 200 local crisis centers, combining custom local care and resources with national standards and best practices. The new, shorter phone number makes it easier for people to remember and access mental health crisis services (though the previous 1 (800) 273-8255 number will continue to function indefinitely) and represents a continued commitment to those in crisis.

*Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-9355(WELL) - 24/7 on call*  
Free and confidential therapy services, workshops, and training for situations related to gender- and power-based harm (including sexual assault, intimate partner violence, and stalking).

*Office for Equity, Equal Opportunity, and Title IX (EEO-TIX) - (213) 740-5086*  
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*  
Avenue to report incidents of bias, hate crimes, and microaggressions to the Office for Equity, Equal Opportunity, and Title for appropriate investigation, supportive measures, and response.

*The Office of Student Accessibility Services (OSAS) - (213) 740-0776*  
OSAS ensures equal access for students with disabilities through providing academic accommodations and auxiliary aids in accordance with federal laws and university policy.

*USC Campus Support and Intervention - (213) 740-0411*  
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

*Diversity, Equity and Inclusion - (213) 740-2101*  
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*  
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) 442-1200 - 24/7 on call*  
Non-emergency assistance or information.

*Office of the Ombuds - (213) 821-9556 (UPC) / (323-442-0382 (HSC)*  
A safe and confidential place to share your USC-related issues with a University Ombuds who will work with you to explore options or paths to manage your concern.

*Occupational Therapy Faculty Practice - (323) 442-2850 or [otfp@med.usc.edu](mailto:otfp@med.usc.edu)*

Confidential Lifestyle Redesign services for USC students to support health promoting habits and routines that enhance quality of life and academic performance.