Psych/CSCI 626: Computational Social Sciences
Text as Data

Morteza Dehghani

Fall 2019

E-mail: mdehghan@usc.edu
Office Hours: Mon 10-12
Office: SGM 607
Web: cssl.usc.edu
Class Hours: Th 2-6pm
Class Room: BCI 266

Course Description

Topics in computational social sciences, or Text as Data, focuses on applications of natural language processing, guided by psychological theories, for identifying various social and cognitive properties evident in human related big data. In this course, we will survey state-of-the-art techniques, and applications of such techniques, 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 social sciences, who are interested in using machine-learning techniques for analysis of data. Also, this course may be of interest to PhD students in communications and the business school.

Learning Objectives

This course is designed to survey current state of research in automated analysis of language within the domain of psychology. It should be noted that the purpose of this class is not to teach text analysis, but to survey how the established methods are used within the social sciences. Optional reading material will be provided for students unfamiliar with the different classes of text analysis.

- **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. Homework assignments will be posted on Blackboard.

Required Readings and Supplementary Materials

- Salganik, M. J. (2017). *Bit by bit: social research in the digital age*. Princeton University Press
- Pennebaker, J. (2011). *The secret life of pronouns: What our words say about us*. New York, NY: Bloomsbury

Description and Assessment of Assignments

1. Paper presentation. Each student will present a set of papers related to one of the topics discussed in class.
2. Reaction paragraphs. Students are asked to write a short note, one or two paragraphs in length, about their reaction to the reading assignments of the week. These can be a quick summary of the material, comments about the subject area, or a critique of a particular theory or experiment. I will read these paragraphs before each class, and will use them to guide the discussion in class.
3. 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

- **15%** Participation
- **20%** Paper Presentations
- **20%** Reaction Paragraphs
- **15%** Final Project Status Updates
- **10%** Final Project Presentation
- **20%** Final Project Write up

Assignment Submission Policy

All assignments are due on Thursdays at 10am. Assignments turned in any later than 10:10am 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.

Schedule and weekly learning goals

The schedule is tentative and subject to change.

Week 01, 08/29: Introduction to Computational Social Sciences 1/2

- Ruths, D. and Pfeffer, J. (2014). Social media for large studies of behavior. *Science*, 346(6213):1063–1064

- Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., et al. (2009). Computational social science. *Science*, 323(5915):721–723

- Yarkoni, T. (2012). Psychoinformatics: New horizons at the interface of the psychological and computing sciences. *Current Directions in Psychological Science*, 21(6):391–397

- King, G. (2011). Ensuring the data-rich future of the social sciences. *science*, 331(6018):719–721

Week 02, 09/05: Introduction to Computational Social Sciences 2/2

- Salganik, M. J. (2017). *Bit by bit: social research in the digital age*. Princeton University Press

  (Chapters 1-3)

- Choose one:
  - Adjerid, I. and Kelley, K. (2018). Big data in psychology: A framework for research advancement. *American Psychologist*, 73(7):899
  - Chen, E. E. and Wojcik, S. P. (2016). A practical guide to big data research in psychology. *Psychological Methods*, 21(4):458

Week 03, 09/12: Dictionary Methods 1/2

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

Week 04, 09/19: Dictionary Methods 2/2

- Back, M. D., Küfner, A. C., and Egloff, B. (2010). The emotional timeline of september 11, 2001. *Psychological Science*, 21(10):1417–1419
- Pury, C. L. (2011). Automation can lead to confounds in text analysis: Back, küfner, and egloff (2010) and the not-so-angry americans. *Psychological science*, 22(6):835
- Back, M. D., Küfner, A. C., and Egloff, B. (2011). "automatic or the people?": Anger on september 11, 2001, and lessons learned for the analysis of large digital data sets. *Psychological Science*, 22(6):837
- Iliev, R., Hoover, J., Dehghani, M., and Axelrod, R. (2016). Linguistic positivity in historical texts reflects dynamic environmental and psychological factors. *Proceedings of the National Academy of Sciences*, 113(49):E7871–E7879
- Choose two:
  - Dehghani, M., Bang, M., Medin, D., Marin, A., Leddon, E., and Waxman, S. (2013). Epistemologies in the text of children's books: Native-and non-native-authored books. *International Journal of Science Education*, 35(13):2133–2151
  - Iliev, R. and Axelrod, R. (2016). Does causality matter more now? increase in the proportion of causal language in english texts. *Psychological science*, 27(5):635–643
  - Frimer, J. A., Aquino, K., Gebauer, J. E., Zhu, L. L., and Oakes, H. (2015). A decline in prosocial language helps explain public disapproval of the us congress. *Proceedings of the National Academy of Sciences*, 112(21):6591–6594

Week 05, 09/26: Differential Language Analysis & Project Proposals

- Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Ramones, S. M., Agrawal, M., Shah, A., Kosinski, M., Stillwell, D., Seligman, M. E., et al. (2013). Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one*, 8(9):e73791
- 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
- Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., Jha, S., Agrawal, M., Dziurzynski, L. A., Sap, M., et al. (2015). Psychological language on twitter predicts county-level heart disease mortality. *Psychological science*, 26(2):159–169
- Brown, N. J. L. and Coyne, J. (2018). No evidence that twitter language reliably predicts heart disease: A reanalysis of eichstaedt et al. (2015a)
- Eichstaedt, j. C., Schwartz, H. A., Giorgi, S., Kern, M. L., Park, G., Sap, M., Labarthe, D. R., Larson, E. E., Seligman, M., and Ungar, L. H. (2018). More evidence that twitter language predicts heart disease: A response and replication

Week 06, 10/03: Latent Semantic Analysis & Distributed Dictionary Representations

- Dam, G. and Kaufmann, S. (2008). Computer assessment of interview data using latent semantic analysis. *Behavior Research Methods*, 40(1):8–20
- Sagi, E. and Dehghani, M. (2014). Measuring moral rhetoric in text. *Social science computer review*, 32(2):132–144
- Dehghani, M., Johnson, K., Hoover, J., Sagi, E., Garten, J., Parmar, N. J., Vaisey, S., Iliev, R., and Graham, J. (2016). Purity homophily in social networks. *Journal of Experimental Psychology: General*, 145(3):366
- Garten, J., Hoover, J., Johnson, K. M., Boghrati, R., Iskiwitch, C., and Dehghani, M. (2018). Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behavior research methods*, 50(1):344–361
- Hoover, J., Johnson, K., Boghrati, R., Graham, J., and Dehghani, M. (2018). Moral framing and charitable donation: Integrating exploratory social media analyses and confirmatory experimentation. *Collabra: Psychology*, 4(1)

Week 07, 10/10: Neural Networks 1/3

- Lin, Y., Hoover, J., Portillo-Wightman, G., Park, C., Dehghani, M., and Ji, H. (2018). Acquiring background knowledge to improve moral value prediction. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 552–559. IEEE
- 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
- Hoover, J., Atari, M., Davani, A. M., Kennedy, B., Portillo-Wightman, G., Yeh, L., Kogon, D., and Dehghani, M. (2019). Bound in hatred: The role of group-based morality in acts of hate
- Bhatia, S., Goodwin, G. P., and Walasek, L. (2018). Trait associations for hillary clinton and donald trump in news media: A computational analysis. *Social Psychological and Personality Science*, 9(2):123–130

Week 08, 10/17: Fall Recess

Week 09, 10/24: Neural Networks 2/3

- Garten, J., Kennedy, B., Hoover, J., Sagae, K., and Dehghani, M. (2019a). Incorporating demographic embeddings into language understanding. *Cognitive science*, 43(1):e12701

- Garten, J., Kennedy, B., Sagae, K., and Dehghani, M. (2019b). Measuring the importance of context when modeling language comprehension. *Behavior research methods*, 51(2):480–492

- Bhatia, S. (2017). Associative judgment and vector space semantics. *Psychological Review*, 124(1):1

- Mostafazadeh Davani, A., Yeh, L., Atari, M., Kennedy, B., Portillo-Wightman, G., Gonzalez, E., Delong, N., Bhatia, R., Mirinjian, A., Ren, X., and Dehghani, M. (2019). Reporting the unreported: Event extraction for analyzing the local representation of hate crimes. In *In the Proceedings of Empirical Methods in Natural Language Processing (EMNLP)*

Week 10, 10/31: Neural Networks 3/3 & Project Update I

- Caliskan, A., Bryson, J. J., and Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186

- 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

- Choose three:

  - Font, J. E. and Costa-Jussa, M. R. (2019). Equalizing gender biases in neural machine translation with word embeddings techniques. *arXiv preprint arXiv:1901.03116*

  - Blodgett, S. L. and O'Connor, B. (2017). Racial disparity in natural language processing: A case study of social media african-american english. *arXiv preprint arXiv:1707.00061*

  - Gonen, H. and Goldberg, Y. (2019). Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. *arXiv preprint arXiv:1903.03862*

  - Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., and Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in neural information processing systems*, pages 4349–4357

Week 11, 11/07: Other methods

- Mehl, M. R., Raison, C. L., Pace, T. W., Arevalo, J. M., and Cole, S. W. (2017). Natural language indicators of differential gene regulation in the human immune system. *Proceedings of the National Academy of Sciences*, 114(47):12554–12559

- 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

- Boghrati, R., Hoover, J., Johnson, K. M., Garten, J., and Dehghani, M. (2018). Conversation level syntax similarity metric. *Behavior research methods*, 50(3):1055–1073
- 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. (2019). How real is moral contagion in online social networks? In *Proceedings of the Cognitive Science Society*

Week 12, 11/14: Clinical & cognitive applications & Project Update II

- Resnik, P., Armstrong, W., Claudino, L., Nguyen, T., Nguyen, V.-A., and Boyd-Graber, J. (2015). Beyond lda: exploring supervised topic modeling for depression-related language in twitter. In *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 99–107
- Walsh, C. G., Ribeiro, J. D., and Franklin, J. C. (2017). Predicting risk of suicide attempts over time through machine learning. *Clinical Psychological Science*, 5(3):457–469
- 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
- Jain, S. and Huth, A. (2018). Incorporating context into language encoding models for fmri. In *Advances in Neural Information Processing Systems*, pages 6628–6637

Week 13, 11/21: Ethics

- Salganik, M. J. (2017). *Bit by bit: social research in the digital age*. Princeton University Press (Chapters 6-7)
- Wienberg, C. and Gordon, A. S. (2015). Insights on privacy and ethics from the web's most prolific storytellers. In *Proceedings of the ACM Web Science Conference*, page 22. ACM
- Hovy, D. and Spruit, S. L. (2016). The social impact of natural language processing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598
- 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 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 Section 11, *Behavior Violating University Standards* `https://scampus.usc.edu/1100-behavior-violating-university-standards-and-appropriate-sanctions/`. Other forms of academic dishonesty are equally unacceptable. See additional information in *SCampus* and university policies on scientific misconduct, `http://policy.usc.edu/scientific-misconduct/`.

Discrimination, sexual assault, and harassment are not tolerated by the university. You are encouraged to report any incidents to the *Office of Equity and Diversity* `http://equity.usc.edu/` or to the *Department of Public Safety* `http://capsnet.usc.edu/department/department-public-safety/online-forms/contact-us`. This is important for the safety whole USC community. Another member of the university community — such as a friend, classmate, advisor, or faculty member — can help initiate the report, or can initiate the report on behalf of another person. *The Center for Women and Men* `http://www.usc.edu/student-affairs/cwm/` provides 24/7 confidential support, and the sexual assault resource center webpage `sarc@usc.edu` describes reporting options and other resources.

Support Systems

A number of USC's schools provide support for students who need help with scholarly writing. Check with your advisor or program staff to find out more. Students whose primary language is not English should check with the *American Language Institute* `http://dornsife.usc.edu/ali`, which sponsors courses and workshops specifically for international graduate students. *The Office of Disability Services and Programs* `http://sait.usc.edu/academicsupport/centerprograms/dsp/home_index.html` provides certification for students with disabilities and helps arrange the relevant accommodations. If an officially declared emergency makes travel to campus infeasible, *USC Emergency Information* `http://emergency.usc.edu/will` provide safety and other updates, including ways in which instruction will be continued by means of blackboard, teleconferencing, and other technology.

References

Adjerid, I. and Kelley, K. (2018). Big data in psychology: A framework for research advancement. *American Psychologist*, 73(7):899.

Back, M. D., Küfner, A. C., and Egloff, B. (2010). The emotional timeline of september 11, 2001. *Psychological Science*, 21(10):1417–1419.

Back, M. D., Küfner, A. C., and Egloff, B. (2011). "automatic or the people?": Anger on september 11, 2001, and lessons learned for the analysis of large digital data sets. *Psychological Science*, 22(6):837.

Bhatia, S. (2017). Associative judgment and vector space semantics. *Psychological Review*, 124(1):1.

Bhatia, S., Goodwin, G. P., and Walasek, L. (2018). Trait associations for hillary clinton and donald trump in news media: A computational analysis. *Social Psychological and Personality Science*, 9(2):123–130.

Blodgett, S. L. and O'Connor, B. (2017). Racial disparity in natural language processing: A case study of social media african-american english. *arXiv preprint arXiv:1707.00061*.

Boghrati, R., Hoover, J., Johnson, K. M., Garten, J., and Dehghani, M. (2018). Conversation level syntax similarity metric. *Behavior research methods*, 50(3):1055–1073.

Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., and Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in neural information processing systems*, pages 4349–4357.

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.

Brown, N. J. L. and Coyne, J. (2018). No evidence that twitter language reliably predicts heart disease: A reanalysis of eichstaedt et al. (2015a).

Burton, J. W., Cruz, N., and Hahn, U. (2019). How real is moral contagion in online social networks? In *Proceedings of the Cognitive Science Society*.

Caliskan, A., Bryson, J. J., and Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.

Chen, E. E. and Wojcik, S. P. (2016). A practical guide to big data research in psychology. *Psychological Methods*, 21(4):458.

Dam, G. and Kaufmann, S. (2008). Computer assessment of interview data using latent semantic analysis. *Behavior Research Methods*, 40(1):8–20.

Dehghani, M., Bang, M., Medin, D., Marin, A., Leddon, E., and Waxman, S. (2013). Epistemologies in the text of children's books: Native-and non-native-authored books. *International Journal of Science Education*, 35(13):2133–2151.

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.

Dehghani, M., Johnson, K., Hoover, J., Sagi, E., Garten, J., Parmar, N. J., Vaisey, S., Iliev, R., and Graham, J. (2016). Purity homophily in social networks. *Journal of Experimental Psychology: General*, 145(3):366.

Eichstaedt, j. C., Schwartz, H. A., Giorgi, S., Kern, M. L., Park, G., Sap, M., Labarthe, D. R., Larson, E. E., Seligman, M., and Ungar, L. H. (2018). More evidence that twitter language predicts heart disease: A response and replication.

Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., Jha, S., Agrawal, M., Dziurzynski, L. A., Sap, M., et al. (2015). Psychological language on twitter predicts county-level heart disease mortality. *Psychological science*, 26(2):159–169.

Font, J. E. and Costa-Jussa, M. R. (2019). Equalizing gender biases in neural machine translation with word embeddings techniques. *arXiv preprint arXiv:1901.03116*.

Frimer, J. A., Aquino, K., Gebauer, J. E., Zhu, L. L., and Oakes, H. (2015). A decline in prosocial language helps explain public disapproval of the us congress. *Proceedings of the National Academy of Sciences*, 112(21):6591–6594.

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.

Garten, J., Hoover, J., Johnson, K. M., Boghrati, R., Iskiwitch, C., and Dehghani, M. (2018). Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behavior research methods*, 50(1):344–361.

Garten, J., Kennedy, B., Hoover, J., Sagae, K., and Dehghani, M. (2019a). Incorporating demographic embeddings into language understanding. *Cognitive science*, 43(1):e12701.

Garten, J., Kennedy, B., Sagae, K., and Dehghani, M. (2019b). Measuring the importance of context when modeling language comprehension. *Behavior research methods*, 51(2):480–492.

Gonen, H. and Goldberg, Y. (2019). Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. *arXiv preprint arXiv:1903.03862*.

Hoover, J., Atari, M., Davani, A. M., Kennedy, B., Portillo-Wightman, G., Yeh, L., Kogon, D., and Dehghani, M. (2019). Bound in hatred: The role of group-based morality in acts of hate.

Hoover, J., Johnson, K., Boghrati, R., Graham, J., and Dehghani, M. (2018). Moral framing and charitable donation: Integrating exploratory social media analyses and confirmatory experimentation. *Collabra: Psychology*, 4(1).

Hovy, D. and Spruit, S. L. (2016). The social impact of natural language processing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598.

Iliev, R. and Axelrod, R. (2016). Does causality matter more now? increase in the proportion of causal language in english texts. *Psychological science*, 27(5):635–643.

Iliev, R., Hoover, J., Dehghani, M., and Axelrod, R. (2016). Linguistic positivity in historical texts reflects dynamic environmental and psychological factors. *Proceedings of the National Academy of Sciences*, 113(49):E7871–E7879.

Jain, S. and Huth, A. (2018). Incorporating context into language encoding models for fmri. In *Advances in Neural Information Processing Systems*, pages 6628–6637.

King, G. (2011). Ensuring the data-rich future of the social sciences. *science*, 331(6018):719–721.

Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., et al. (2009). Computational social science. *Science*, 323(5915):721–723.

Lin, Y., Hoover, J., Portillo-Wightman, G., Park, C., Dehghani, M., and Ji, H. (2018). Acquiring background knowledge to improve moral value prediction. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 552–559. IEEE.

Mehl, M. R., Raison, C. L., Pace, T. W., Arevalo, J. M., and Cole, S. W. (2017). Natural language indicators of differential gene regulation in the human immune system. *Proceedings of the National Academy of Sciences*, 114(47):12554–12559.

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.

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.

Mostafazadeh Davani, A., Yeh, L., Atari, M., Kennedy, B., Portillo-Wightman, G., Gonzalez, E., Delong, N., Bhatia, R., Mirinjian, A., Ren, X., and Dehghani, M. (2019). Reporting the unreported: Event extraction for analyzing the local representation of hate crimes. In *In the Proceedings of Empirical Methods in Natural Language Processing (EMNLP)*.

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.

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

Pury, C. L. (2011). Automation can lead to confounds in text analysis: Back, küfner, and egloff (2010) and the not-so-angry americans. *Psychological science*, 22(6):835.

Resnik, P., Armstrong, W., Claudino, L., Nguyen, T., Nguyen, V.-A., and Boyd-Graber, J. (2015). Beyond lda: exploring supervised topic modeling for depression-related language in twitter. In *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 99–107.

Ruths, D. and Pfeffer, J. (2014). Social media for large studies of behavior. *Science*, 346(6213):1063–1064.

Sagi, E. and Dehghani, M. (2014). Measuring moral rhetoric in text. *Social science computer review*, 32(2):132–144.

Salganik, M. J. (2017). *Bit by bit: social research in the digital age*. Princeton University Press.

Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Ramones, S. M., Agrawal, M., Shah, A., Kosinski, M., Stillwell, D., Seligman, M. E., et al. (2013). Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one*, 8(9):e73791.

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.

Walsh, C. G., Ribeiro, J. D., and Franklin, J. C. (2017). Predicting risk of suicide attempts over time through machine learning. *Clinical Psychological Science*, 5(3):457–469.

Wienberg, C. and Gordon, A. S. (2015). Insights on privacy and ethics from the web's most prolific storytellers. In *Proceedings of the ACM Web Science Conference*, page 22. ACM.

Yarkoni, T. (2012). Psychoinformatics: New horizons at the interface of the psychological and computing sciences. *Current Directions in Psychological Science*, 21(6):391–397.