

References

- [1] L. Adamic and E. Adar. Friends and neighbors on the web. *Social Networks*, 25, 2001.
- [2] Y. Bachrach, M. Kosinski, T. Graepel, P. Kohli, and D. Stillwell. Personality and patterns of Facebook usage. In *Proceedings of ACM WebSci*, pages 24–32, 2012.
- [3] L. Backstrom and J. Kleinberg. Romantic partnerships and the dispersion of social ties: A network analysis of relationship status on Facebook. In *Proceedings of CSCW*, 2014.
- [4] S. Bergsma, M. Dredze, B. Van Durme, T. Wilson, and D. Yarowsky. Broadly improving user classification via communication-based name and location clustering on Twitter. In *Proceedings of NAACL-HLT*, pages 1010–1019, 2013.
- [5] S. Bergsma, M. Post, and D. Yarowsky. Stylometric analysis of scientific articles. In *Proceedings of NAACL-HLT*, pages 327–337, 2012.
- [6] S. Bergsma and B. Van Durme. Using conceptual class attributes to characterize social media users. In *Proceedings of ACL*, pages 710–720, 2013.
- [7] A. Bifet and E. Frank. Sentiment knowledge discovery in Twitter streaming data. In *Proceedings of the 13th International Conference on Discovery Science*, DS’10, pages 1–15, 2010.
- [8] A. Bifet, G. Holmes, B. Pfahringer, and R. Gavaldà. Detecting sentiment change in Twitter streaming data. *Journal of Machine Learning Research*, 17:5–11, 2011.
- [9] A. Bifet, G. Holmes, B. Pfahringer, P. Kranen, H. Kremer, T. Jansen, and T. Seidl. MOA: Massive online analysis, a framework for stream classification and clustering. *Journal of Machine Learning Research*, 11:44–50, 2010.
- [10] M. Bilgic and L. Getoor. Active inference for collective classification. In *Proceedings of AAAI*, pages 1652–1655, 2010.
- [11] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., 2006.
- [12] J. Bollen, B. Gonçalves, G. Ruan, and H. Mao. Happiness is assortative in online social networks. *Artificial life*, 17(3):237–251, 2011.
- [13] J. Bollen, H. Mao, and A. Pepe. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. In *Proceedings of ICWSM*, 2011.
- [14] J. Bollen, H. Mao, and X. Zeng. Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1):1–8, 2011.
- [15] J. D. Burger, J. Henderson, G. Kim, and G. Zarrella. Discriminating gender on Twitter. In *Proceedings of EMNLP*, pages 1301–1309, 2011.
- [16] J. D. Burger and J. C. Henderson. An exploration of observable features related to blogger age. In *AAAI Spring Symposium: Computational Approaches to Analyzing Weblogs*, pages 15–20, 2006.

- [17] C. Callison-Burch. Fast, cheap, and creative: Evaluating translation quality using Amazon’s Mechanical Turk. In *Proceedings of EMNLP*, pages 286–295, 2009.
- [18] M. Cha, H. Haddadi, F. Benevenuto, and K. P. Gummadi. Measuring user influence in Twitter: The million follower fallacy. In *Proceedings of ICWSM*, 2010.
- [19] N. Cheng, R. Chandramouli, and K. Subbalakshmi. Author gender identification from text. *Digital Investigation*, 8(1):78–88, 2011.
- [20] M. Ciot, M. Sonderegger, and D. Ruths. Gender inference of Twitter users in non-english contexts. In *Proceedings of EMNLP*, pages 1136–1145, 2013.
- [21] R. Cohen and D. Ruths. Classifying political orientation on Twitter: It’s not easy! In *Proceedings of ICWSM*, 2013.
- [22] M. D. Conover, B. Gonçalves, J. Ratkiewicz, A. Flammini, and F. Menczer. Predicting the political alignment of Twitter users. In *Proceedings of Social Computing*, 2011.
- [23] M. D. Conover, J. Ratkiewicz, M. Francisco, B. Gonc, A. Flammini, and F. Menczer. Political polarization on Twitter. In *Proceedings of ICWSM*, pages 89–96, 2011.
- [24] G. Coppersmith, M. Dredze, and C. Harman. Quantifying mental health signals in twitter. In *ACL Workshop on Computational Linguistics and Clinical Psychology*, 2014.
- [25] G. Coppersmith, C. Harman, and M. Dredze. Measuring post traumatic stress disorder in Twitter. In *Proceedings of ICWSM*, 2014.
- [26] G. A. Coppersmith and C. E. Priebe. Vertex nomination via content and context. *Computing Research Repository*, 2012.
- [27] L. Coviello, Y. Sohn, A. D. Kramer, C. Marlow, M. Franceschetti, N. A. Christakis, and J. H. Fowler. Detecting emotional contagion in massive social networks. *PloS one*, 9(3):e90315, 2014.
- [28] M. De Choudhury, S. Counts, and E. Horvitz. Predicting postpartum changes in emotion and behavior via social media. In *Proceedings of SIGCHI*, pages 3267–3276, 2013.
- [29] M. De Choudhury and S. De. Mental health discourse on reddit: Self-disclosure, social support, and anonymity. 2014.
- [30] M. De Choudhury, M. Gamon, and S. Counts. Happy, nervous or surprised? classification of human affective states in social media. In *Proceedings of ICWSM*, 2012.
- [31] M. De Choudhury, M. Gamon, S. Counts, and E. Horvitz. Predicting depression via social media. In *Proceedings of ICWSM*, 2013.
- [32] D. Ediger, K. Jiang, J. Riedy, D. A. Bader, and C. Corley. Massive social network analysis: Mining twitter for social good. In *Proceedings of ICPP*, 2010.
- [33] J. Eisenstein, B. O’Connor, N. A. Smith, and E. P. Xing. A latent variable model for geographic lexical variation. In *Proceedings of EMNLP*, pages 1277–1287, 2010.
- [34] J. Eisenstein, B. O’Connor, N. A. Smith, and E. P. Xing. Diffusion of lexical change in social media. *PloS one*, 9(11):e113114, 2014.

- [35] P. Ekman. An argument for basic emotions. *Cognition & Emotion*, 6(3-4):169–200, May 1992.
- [36] A. A. A. Esmin, R. de Oliveira, and S. Matwin. Hierarchical classification approach to emotion recognition in Twitter. In *Proceedings of ICMLA*, volume 2, pages 381–385, 2012.
- [37] R. Fan, J. Zhao, Y. Chen, and K. Xu. Anger is more influential than joy: sentiment correlation in Weibo. *arXiv preprint arXiv:1309.2402*, 2013.
- [38] R. E. Fan, K. W. Chang, C. J. Hsieh, X. R. Wang, and C. J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008.
- [39] Q. Fang, J. Sang, C. Xu, and M. Hossain. Exploiting interaction with multimedia information for relational user attribute inference. 2015.
- [40] K. Filippova. User demographics and language in an implicit social network. In *Proceedings of EMNLP-CoNLL*, 2012.
- [41] N. Garera and D. Yarowsky. Modeling latent biographic attributes in conversational genres. In *Proceedings of ACL*, pages 710–718, 2009.
- [42] N. Garera and D. Yarowsky. Structural, transitive and latent models for biographic fact extraction. In *Proceedings of EACL*, pages 300–308, 2009.
- [43] D. Gayo-Avello. No, you cannot predict elections with Twitter. *Internet Computing, IEEE*, 16(6):91–94, 2012.
- [44] D. Gayo-Avello, P. T. Metaxas, and E. Mustafaraj. Limits of Electoral Predictions Using Twitter. In *Proceedings of ICWSM*, 2011.
- [45] A. Go, R. Bhayani, and L. Huang. Twitter Sentiment Classification using Distant Supervision. Technical report, 2009.
- [46] J. Golbeck and D. Hansen. Computing political preference among Twitter followers. In *Proceedings of the SIGCHI*, pages 1105–1108, 2011.
- [47] J. Golbeck, C. Robles, M. Edmondson, and K. Turner. Predicting personality from Twitter. In *Proceedings of SocialCom/PASSAT*, 2011.
- [48] R. González-Ibáñez, S. Muresan, and N. Wacholder. Identifying sarcasm in Twitter: A closer look. In *Proceedings of ACL*, pages 581–586, 2011.
- [49] S. Goswami, S. Sarkar, and M. Rustagi. Stylometric analysis of bloggers’ age and gender. In *Proceedings of CWSM*, 2009.
- [50] A. Goyal, H. Daumé, III, and S. Venkatasubramanian. Streaming for large scale NLP: language modeling. In *Proceedings of NAACL*, pages 512–520, 2009.
- [51] Y. H. Hassan Saif, Miriam Fernandez and H. Alani. Evaluation datasets for Twitter sentiment analysis: A survey and a new dataset, the sts-gold. *First ESSEM workshop*, 2013.
- [52] E. Hatfield and J. T. Cacioppo. *Emotional contagion*. Cambridge university press, 1994.

- [53] D. Hovy. Demographic factors improve classification performance. *Proceedings of ACL*, 2015.
- [54] D. Kahneman and A. Deaton. High income improves evaluation of life but not emotional well-being. *Proceedings of the National Academy of Sciences*, 107(38):16489–16493, 2010.
- [55] M. L. Kern, J. C. Eichstaedt, H. A. Schwartz, L. Dziurzynski, L. H. Ungar, D. J. Stillwell, M. Kosinski, S. M. Ramones, and M. E. Seligman. The online social self an open vocabulary approach to personality. *Assessment*, 21(2):158–169, 2014.
- [56] M. L. Kern, J. C. Eichstaedt, H. A. Schwartz, G. Park, L. H. Ungar, D. J. Stillwell, M. Kosinski, L. Dziurzynski, and M. E. Seligman. From sooo excited!!! to so proud: Using language to study development. *Developmental psychology*, 50(1):178, 2014.
- [57] E. Kharitonov and P. Serdyukov. Gender-aware re-ranking. In *SIGIR*, pages 1081–1082, 2012.
- [58] S. Kim, J. Bak, and A. H. Oh. Do you feel what I feel? Social aspects of emotions in Twitter conversations. In *Proceedings of ICWSM*, 2012.
- [59] F. Kivran-Swaine, S. Brody, N. Diakopoulos, and M. Naaman. Of joy and gender: emotional expression in online social networks. In *Proceedings of CSCWC*, pages 139–142, 2012.
- [60] A. Kokkos and T. Tzouramanis. A robust gender inference model for online social networks and its application to linkedin and twitter. *First Monday*, 19(9), 2014.
- [61] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. 2009.
- [62] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *National Academy of Sciences*, 2013.
- [63] H. Kwak, C. Lee, H. Park, and S. Moon. What is Twitter, a social network or a news media? In *Proceedings of WWW*, pages 591–600, 2010.
- [64] V. Lampos. On voting intentions inference from twitter content: a case study on uk 2010 general election. *Computing Research Repository*, 2012.
- [65] V. Lampos, T. De Bie, and N. Cristianini. Flu detector: tracking epidemics on Twitter. In *Proceedings of ECML PKDD*, pages 599–602, 2010.
- [66] V. Lampos, D. Preotiuc-Pietro, and T. Cohn. A user-centric model of voting intention from social media. In *Proceedings of ACL*, pages 993–1003, 2013.
- [67] F. Laws. *Effective Active Learning for Complex Natural Language Processing Tasks*. PhD thesis, University of Stuttgart, 2012.
- [68] K. Lewis, J. Kaufman, M. Gonzalez, A. Wimmer, and N. Christakis. Tastes, ties, and time: A new social network dataset using facebook.com. *Social Networks*, 30(4):330–342, 2008.
- [69] W. Liu and D. Ruths. What’s in a name? using first names as features for gender inference in twitter. In *AAAI Spring Symposium: Analyzing Microtext*, 2013.

- [70] L. Lovasz. Random walks on graphs: A survey, 1993.
- [71] I. Mackinnon. Age and geographic inferences of the livejournal social network. In *Statistical Network Analysis Workshop*, 2006.
- [72] J. Marquardt, G. Farnadi, G. Vasudevan, M.-F. Moens, S. Davalos, A. Teredesai, and M. De Cock. Age and gender identification in social media. *Proceedings of CLEF 2014 Evaluation Labs*, 2014.
- [73] D. Maynard and A. Funk. Automatic detection of political opinions in tweets. In *Proceedings of ESWC*, pages 88–99, 2012.
- [74] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: homophily in social networks. *Annual Review of Sociology*, 27(1):415–444, 2001.
- [75] A. Mislove, S. Lehmann, Y.-Y. Ahn, J.-P. Onnela, and J. N. Rosenquis. Understanding the demographics of Twitter users. In *Proceedings of ICWSM*, pages 554–557, 2011.
- [76] S. Mohammad. #emotional tweets. In *Proceedings of SemEval 2012*, pages 246–255, 2012.
- [77] S. Mohammad and S. Kiritchenko. Using nuances of emotion to identify personality. In *Proceedings of ICWSM*, 2013.
- [78] S. Mohammad and T. Yang. Tracking sentiment in mail: How genders differ on emotional axes. In *Proceedings of WASSA*, pages 70–79, 2011.
- [79] S. M. Mohammad and S. Kiritchenko. Using hashtags to capture fine emotion categories from tweets. *Computational Intelligence*, 2014.
- [80] S. M. Mohammad, S. Kiritchenko, and X. Zhu. Nrc-canada: Building the state-of-the-art in sentiment analysis of tweets. In *Proceedings of SemEval*, June 2013.
- [81] P. Nakov, S. Rosenthal, Z. Kozareva, V. Stoyanov, A. Ritter, and T. Wilson. Semeval-2013 task 2: Sentiment analysis in Twitter. In *Proceedings of SemEval*, pages 312–320, 2013.
- [82] D. Nguyen, R. Gravel, D. Trieschnigg, and T. Meder. "How old do you think I am?" A study of language and age in Twitter. In *Proceedings of ICWSM*, pages 439–448, 2013.
- [83] D. Nguyen, N. A. Smith, and C. P. Rosé. Author age prediction from text using linear regression. In *Proceedings of LaTeCH*, pages 115–123, 2011.
- [84] B. O’Connor, R. Balasubramanyan, B. R. Routledge, and N. A. Smith. From tweets to polls: Linking text sentiment to public opinion time series. In *Proceedings of ICWSM*, pages 122–129, 2010.
- [85] B. O’Connor, J. Eisenstein, E. P. Xing, and N. A. Smith. A mixture model of demographic lexical variation. In *Proceedings of NIPS MLSC*, 2010.
- [86] F. Olsson. A literature survey of active machine learning in the context of natural language processing. Technical Report 06, 2009.

- [87] M. Ott, Y. Choi, C. Cardie, and J. T. Hancock. Finding deceptive opinion spam by any stretch of the imagination. In *Proceedings of ACL*, pages 309–319, 2011.
- [88] O. Owoputi, B. O’Connor, C. Dyer, K. Gimpel, N. Schneider, and N. A. Smith. Improved part-of-speech tagging for online conversational text with word clusters. In *Proceedings of NAACL*, pages 380–390, 2013.
- [89] A. Pak and P. Paroubek. Twitter as a corpus for sentiment analysis and opinion mining. In *LREC*, 2010.
- [90] B. Pang and L. Lee. Opinion mining and sentiment analysis. *Foundations of Trends in IR*, 2(1-2):1–135, Jan. 2008.
- [91] B. Pang, L. Lee, and S. Vaithyanathan. Thumbs up?: sentiment classification using machine learning techniques. In *Proceedings of EMNLP*, pages 79–86, 2002.
- [92] M. J. Paul and M. Dredze. You Are What You Tweet: Analyzing Twitter for Public Health. *Artificial Intelligence*, pages 265–272, 2011.
- [93] L. Pearl and M. Steyvers. Identifying emotions, intentions, and attitudes in text using a game with a purpose. In *Proceedings of the Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 71–79, 2010.
- [94] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and d. Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [95] C. Peersman, W. Daelemans, and L. Van Vaerenbergh. Predicting age and gender in online social networks. In *Proceedings of SMUC*, pages 37–44, 2011.
- [96] M. Pennacchiotti and A.-M. Popescu. Democrats, republicans and starbucks aficionados: user classification in Twitter. In *Proceedings of KDD*, pages 430–438, 2011.
- [97] M. Pennacchiotti and A. M. Popescu. A machine learning approach to Twitter user classification. In *Proceedings of ICWSM*, pages 281–288, 2011.
- [98] J. W. Pennebaker, J. Both, Roger, and M. E. Francis. Linguistic Inquiry and Word Count (LIWC): A Computerized Text Analysis Program, 2001.
- [99] D. Preotiuc-Pietro, J. Eichstaedt, G. Park, M. Sap, L. Smith, V. Tobolsky, H. A. Schwartz, and L. Ungar. The role of personality, age and gender in tweeting about mental illnesses. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality, NAACL*, 2015.
- [100] D. Preotiuc-Pietro, V. Lampos, and N. Aletras. An analysis of the user occupational class through Twitter content. *Proceedings of ACL*, 2015.
- [101] M. Purver and S. Battersby. Experimenting with distant supervision for emotion classification. In *Proceedings of ACL*, pages 482–491, 2012.
- [102] A. Qadir and E. Riloff. Bootstrapped learning of emotion hashtags# hashtags4you. *WASSA 2013*, page 2, 2013.

- [103] D. Rao, M. Paul, C. Fink, D. Yarowsky, T. Oates, and G. Coppersmith. Hierarchical Bayesian models for latent attribute detection in social media. In *Proceedings of ICWSM*, 2011.
- [104] D. Rao, D. Yarowsky, A. Shreevats, and M. Gupta. Classifying latent user attributes in Twitter. In *Proceedings of SMUC*, pages 37–44, 2010.
- [105] P. Resnik. Getting real(-time) with live polling, 2013. <http://vimeo.com/68210812>.
- [106] K. Roberts, M. A. Roach, J. Johnson, J. Guthrie, and S. M. Harabagiu. Empatweet: Annotating and detecting emotions on Twitter. In *Proceedings of LREC*, 2012.
- [107] S. Rosenthal and K. McKeown. Age prediction in blogs: A study of style, content, and online behavior in pre-and post-social media generations. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 763–772. Association for Computational Linguistics, 2011.
- [108] J. Schler, M. Koppel, S. Argamon, and J. Pennebaker. Effects of Age and Gender on Blogging. In *Proceedings of AAAI Spring Symposium on Computational Approaches for Analyzing Weblogs*, 2006.
- [109] H. A. Schwartz, J. C. Eichstaedt, L. Dziurzynski, M. L. Kern, E. Blanco, M. Kosinski, D. Stillwell, M. E. Seligman, and L. H. Ungar. Toward personality insights from language exploration in social media. In *AAAI Spring Symposium: Analyzing Microtext*, 2013.
- [110] H. A. Schwartz, J. C. Eichstaedt, M. L. Kern, L. Dziurzynski, R. E. Lucas, M. Agrawal, G. J. Park, S. K. Lakshmikanth, S. Jha, M. E. Seligman, et al. Characterizing geographic variation in well-being using tweets. In *ICWSM*, 2013.
- [111] H. A. Schwartz, J. C. Eichstaedt, M. L. Kern, L. Dziurzynski, S. M. Ramones, M. Agrawal, A. Shah, M. Kosinski, D. Stillwell, M. E. Seligman, et al. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one*, 8(9):e73791, 2013.
- [112] B. Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison, 2010.
- [113] B. Settles. *Active Learning*. Morgan and Claypool, 2012.
- [114] B. Settles, M. Craven, and S. Ray. Multiple-instance active learning. In *In Advances in Neural Information Processing Systems (NIPS)*, pages 1289–1296, 2008.
- [115] N. A. Smith. Log-linear models, 2004.
- [116] R. Snow, B. O’Connor, D. Jurafsky, and A. Y. Ng. Cheap and fast—but is it good?: Evaluating non-expert annotations for natural language tasks. In *Proceedings of EMNLP*, pages 254–263, 2008.
- [117] C. Strapparava and R. Mihalcea. Semeval-2007 task 14: Affective text. In *Proceedings of SemEval*, pages 70–74, 2007.
- [118] M. Thelwall. Homophily in MySpace. *Journal of the American Society for Information Science and Technology*, 60(2):219–231, 2009.

- [119] R. Valitutti. Wordnet-affect: an affective extension of wordnet. In *Proceedings of LREC*, pages 1083–1086, 2004.
- [120] B. Van Durme. Jerboa: A toolkit for randomized and streaming algorithms. Technical report, Human Language Technology Center of Excellence, 2012.
- [121] B. Van Durme. Streaming analysis of discourse participants. In *Proceedings of EMNLP*, pages 48–58, 2012.
- [122] S. Volkova, G. Coppersmith, and B. Van Durme. Inferring user political preferences from streaming communications. In *Proceedings of ACL*, pages 186–196, 2014.
- [123] S. Volkova and B. Van Durme. Online bayesian models for personal analytics in social media. In *Proceedings of AAAI*, January 2015.
- [124] S. Volkova, T. Wilson, and D. Yarowsky. Exploring demographic language variations to improve multilingual sentiment analysis in social media. In *Proceedings of EMNLP*, 2013.
- [125] S. Volkova and D. Yarowsky. Improving gender prediction of social media users via weighted annotator rationales. In *NIPS 2014 Workshop on Personalization*, December 2014.
- [126] W. Wang, L. Chen, K. Thirunarayan, and A. P. Sheth. Harnessing Twitter “big data” for automatic emotion identification. In *Proceedings of SocialCom*, pages 587–592, 2012.
- [127] N. Weil. Privacy group files FTC complaint over Facebook’s ‘emotional contagion’ study. *PCWorld*, 2014.
- [128] S. Wu, J. M. Hofman, W. A. Mason, and D. J. Watts. Who says what to whom on Twitter. In *Proceedings of WWW*, 2011.
- [129] Y. Yang and J. Eisenstein. A log-linear model for unsupervised text normalization. In *Proceedings of EMNLP*, pages 61–72, 2013.
- [130] A. Yessenalina, Y. Choi, and C. Cardie. Automatically generating annotator rationales to improve sentiment classification. In *Proceedings of ACL*, pages 336–341, 2010.
- [131] W. Youyou, M. Kosinski, and D. Stillwell. Computer-based personality judgments are more accurate than those made by humans. *PNAS*, 2015.
- [132] O. Zaidan, J. Eisner, and C. D. Piatko. Using “annotator rationales” to improve machine learning for text categorization. In *HLT-NAACL*, pages 260–267, 2007.
- [133] O. F. Zaidan and J. Eisner. Modeling annotators: A generative approach to learning from annotator rationales. In *Proceedings of EMNLP 2008*, pages 31–40, October 2008.
- [134] O. F. Zaidan, J. Eisner, and C. Piatko. Machine learning with annotator rationales to reduce annotation cost. In *Proceedings of the NIPS Workshop on Cost Sensitive Learning*, December 2008.
- [135] F. A. Zamal, W. Liu, and D. Ruths. Homophily and latent attribute inference: Inferring latent attributes of Twitter users from neighbors. In *Proceedings of ICWSM*, 2012.
- [136] X. Zhu, S. Kiritchenko, and S. M. Mohammad. NRC-Canada-2014: Recent improvements in the sentiment analysis of tweets. *SemEval*, 2014.