

Acknowledgments. This work was partly supported by Collaborative Research on Science and Society, Swiss Data Science Center, Swiss National Science Foundation (grant 200021_185043), Facebook, Google, Microsoft. We thank Tiziano Piccardi, Lorenzo Tarantino, Dario Pavlo, Seth Vanderwilt for early help, and Spinn3r and Jure Leskovec for facilitating data access.

REFERENCES

- [1] Eugene Agichtein and Luis Gravano. 2000. *Snowball: Extracting Relations From Large Plain-text Collections*. In *Proceedings of the Fifth ACM Conference on Digital Libraries*. <https://doi.org/10.1145/336597.336644>
- [2] Mariana S. C. Almeida, Miguel B. Almeida, and André F. T. Martins. 2014. A Joint Model for Quotation Attribution and Coreference Resolution. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, EACL*. <https://doi.org/10.3115/v1/e14-1005>
- [3] Kevin Burton, Niels Kasch, and Ian Soboroff. 2011. The ICWSM 2011 Spinn3r Dataset. In *Proceedings of the Fifth Annual Conference on Weblogs and Social Media, ICWSM*.
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*. <https://doi.org/10.18653/v1/n19-1423>
- [5] David K. Elson and Kathleen R. McKeown. 2010. Automatic Attribution of Quoted Speech in Literary Narrative. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI*. <http://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1945>
- [6] Marti A. Hearst. 1992. Automatic Acquisition of Hyponyms from Large Text Corpora. In *14th International Conference on Computational Linguistics, COLING*. <https://www.aclweb.org/anthology/C92-2082/>
- [7] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR*. <http://arxiv.org/abs/1412.6980>
- [8] Mark Leibovich. 2019. How Lindsey Graham Went From Trump Skeptic to Trump Sidekick. *The New York Times* (February 25, 2019). <https://www.nytimes.com/2019/02/25/magazine/lindsey-graham-what-happened-trump.html>
- [9] Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP Natural Language Processing Toolkit. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL*, 55–60. <https://doi.org/10.3115/v1/p14-5010>
- [10] Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics, ACL, and the 4th International Joint Conference on Natural Language Processing of the AFNLP*. <https://www.aclweb.org/anthology/P09-1113/>
- [11] Grace Muzny, Michael Fang, Angel X. Chang, and Dan Jurafsky. 2017. A Two-stage Sieve Approach for Quote Attribution. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL*. <https://doi.org/10.18653/v1/e17-1044>
- [12] Chris Newell, Tim Cowlshaw, and David Man. 2018. Quote Extraction and Analysis for News. In *KDD Workshop on Data Science, Journalism & Media, DSJM*.
- [13] Vlad Niculae, Caroline Suen, Justine Zhang, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. 2015. QUOTUS: The Structure of Political Media Coverage as Revealed by Quoting Patterns. In *Proceedings of the 24th International Conference on World Wide Web, WWW*. <https://doi.org/10.1145/2736277.2741688>
- [14] Timothy O’Keefe, Silvia Pareti, James R. Curran, Irena Koprinska, and Matthew Honnibal. 2012. A Sequence Labelling Approach to Quote Attribution. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, EMNLP-CoNLL*. <https://www.aclweb.org/anthology/D12-1072/>
- [15] Timothy O’Keefe, Kellie Webster, James R. Curran, and Irena Koprinska. 2013. Examining the Impact of Coreference Resolution on Quote Attribution. In *Proceedings of the Australasian Language Technology Association Workshop, ALTA*. <https://www.aclweb.org/anthology/U13-1007/>
- [16] Silvia Pareti, Timothy O’Keefe, Ioannis Konstas, James R. Curran, and Irena Koprinska. 2013. Automatically Detecting and Attributing Indirect Quotations. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP*. <https://www.aclweb.org/anthology/D13-1101/>
- [17] Dario Pavlo, Tiziano Piccardi, and Robert West. 2018. Quototrap: Scalable Unsupervised Extraction of Quotation-Speaker Pairs from Large News Corpora via Bootstrapping. In *Proceedings of the Twelfth International Conference on Web and Social Media, ICWSM*. <https://aaai.org/ocs/index.php/ICWSM/ICWSM18/paper/view/17827>
- [18] Bruno Pouliquen, Ralf Steinberger, and Clive Best. 2007. Automatic Detection of Quotations in Multilingual News. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing, RANLP*.
- [19] Andrew Salway, Paul Meurer, Knut Hofland, and Øystein Reigem. 2017. In *Proceedings of the 21st Nordic Conference on Computational Linguistics, NODALIDA*. <http://www.ep.liu.se/ecp/article.asp?issue=131&article=041&volume=>
- [20] Andreas Spitz and Michael Gertz. 2018. Exploring Entity-centric Networks in Entangled News Streams. In *Companion of the The Web Conference, WWW*. <https://doi.org/10.1145/3184558.3188726>
- [21] Caroline Suen, Sandy Huang, Chantat Eksombatchai, Rok Soscic, and Jure Leskovec. 2013. NIFTY: A System for Large Scale Information Flow Tracking and Clustering. In *22nd International World Wide Web Conference, WWW*. <https://doi.org/10.1145/2488388.2488496>
- [22] Thomas Pellissier Tanon, Denny Vrandečić, Sebastian Schaffert, Thomas Steiner, and Lydia Pintscher. 2016. From Freebase to Wikidata: The Great Migration. In *Proceedings of the 25th International Conference on World Wide Web, WWW*. <https://doi.org/10.1145/2872427.2874809>
- [23] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems, NeurIPS*. <http://papers.nips.cc/paper/7181-attention-is-all-you-need>
- [24] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *CoRR* abs/1910.03771 (2019). arXiv:1910.03771 <http://arxiv.org/abs/1910.03771>
- [25] Yi Zhang, Zachary G. Ives, and Dan Roth. 2020. "Who said it, and Why?" Provenance for Natural Language Claims. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL*. <https://www.aclweb.org/anthology/2020.acl-main.406/>