

by the labeling of each claim (i.e. keeping only the claims labeled as true). It is interesting to see which systems are adjudged to have provided the correct answers.

This visualization is available as an interactive demo.²

6. DISCUSSION

We have presented a list of desiderata for a good trustworthiness visualization, and argued that these cannot be satisfied using only simple statistics on the graph. To meet these desiderata, we have shown three different visualization techniques, and commented on their effectiveness.

The first method is using a simple standard bipartite layout of a graph. We argued that this method is infeasible for the purpose because it does not scale elegantly. The second method is using a physics-based layout of the trustworthiness graph. We found that this visualization is useful in discovering large-scale communities of nodes, and for getting a global picture of the data. Unlike the first method, this scales to a large number of nodes. But while it does well at big-picture understanding, it fails to reveal more fine-grained facets. Our final method, based on the adjacency matrix representation of the graph, gives much more control by allowing exploration of the data by different attributes of claims and sources.

7. ACKNOWLEDGEMENTS

We would like to thank the reviewers for their thoughtful and helpful comments and the creators of the d3 library³ for making such an effective and beautiful tool.

This research was supported by Army Research Laboratory (ARL) under agreement W911NF-09-2-0053, by DARPA under agreement number FA8750-13-2-0008, and by the Multimodal Information Access & Synthesis Center at UIUC, part of CCICADA, a DHS Science and Technology Center of Excellence. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. Any opinions, findings, conclusions or recommendations are those of the authors and do not necessarily reflect the view of ARL, DARPA or the U.S. Government.

8. REFERENCES

- [1] Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems*, 30(1):107–117, 1998.
- [2] Jon M Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632, 1999.
- [3] Jeff Pasternack and Dan Roth. Knowing what to believe (when you already know something). In *Proceedings of the 23rd International Conference on Computational Linguistics*, pages 877–885. Association for Computational Linguistics, 2010.
- [4] Xin Luna Dong, Evgeniy Gabrilovich, Kevin Murphy, Van Dang, Wilko Horn, Camillo Lugaresi, Shaohua Sun, and Wei Zhang. Knowledge-Based Trust: Estimating the Trustworthiness of Web Sources. *arXiv.org*, February 2015.
- [5] Xin Luna Dong, Evgeniy Gabrilovich, Jeremy Heitz, Wilko Horn, Kevin Murphy, Shaohua Sun, and Wei Zhang. From data fusion to knowledge fusion. *Proceedings of the VLDB Endowment*, 7(10), 2014.
- [6] David R Karger, Sewoong Oh, and Devavrat Shah. Iterative learning for reliable crowdsourcing systems. In *Advances in neural information processing systems*, pages 1953–1961, 2011.
- [7] Sepandar D Kamvar, Mario T Schlosser, and Hector Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of the 12th international conference on World Wide Web*, pages 640–651. ACM, 2003.
- [8] Jennifer Neville and David Jensen. Iterative classification in relational data. In *Proc. AAAI-2000 Workshop on Learning Statistical Models from Relational Data*, pages 13–20, 2000.
- [9] Dian Yu, Hongzhao Huang, Taylor Cassidy, Heng Ji, Chi Wang, Shi Zhi, Jiawei Han, Clare Voss, and Malik Magdon-Ismael. The wisdom of minority: Unsupervised slot filling validation based on multi-dimensional truth-finding. In *Proc. The 25th International Conference on Computational Linguistics (COLING2014)*, 2014.
- [10] Yifan Hu and Lei Shi. Visualizing large graphs. *Wiley Interdisciplinary Reviews: Computational Statistics*, 7(2):115–136, 2015.
- [11] Helen C. Purchase, Robert F. Cohen, and Murray I James. An experimental study of the basis for graph drawing algorithms. *Journal of Experimental Algorithmics (JEA)*, 2:4, 1997.
- [12] Mohammad Ghoniem, J Fekete, and Philippe Castagliola. A comparison of the readability of graphs using node-link and matrix-based representations. In *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on*, pages 17–24. IEEE, 2004.
- [13] Niklas Elmqvist, Thanh-Nghi Do, Howard Goodell, Nathalie Henry, and J Fekete. Zame: Interactive large-scale graph visualization. In *Visualization Symposium, 2008. PacificVIS'08. IEEE Pacific*, pages 215–222. IEEE, 2008.
- [14] J Fekete. Visualizing networks using adjacency matrices: Progresses and challenges. In *Computer-Aided Design and Computer Graphics, 2009. CAD/Graphics' 09. 11th IEEE International Conference on*, pages 636–638. IEEE, 2009.
- [15] Mihai Surdeanu. Overview of the tac2013 knowledge base population evaluation: English slot filling and temporal slot filling. In *Proceedings of the Sixth Text Analysis Conference (TAC 2013)*, 2013.
- [16] Peter Eades and NicholasC. Wormald. Edge crossings in drawings of bipartite graphs. *Algorithmica*, 11(4):379–403, 1994.
- [17] Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. Gephi: An open source software for exploring and manipulating networks. In *International AAAI Conference on Weblogs and Social Media*, 2009.
- [18] Shawn Martin, W Michael Brown, Richard Klavans, and Kevin W Boyack. Openord: An open-source toolbox for large graph layout. In *IS&T/SPIE Electronic Imaging*, pages 786806–786806. International Society for Optics and Photonics, 2011.

²<http://cogcomp.cs.illinois.edu/~mayhew2/kbpvis/>

³<http://d3js.org/>