PhD proposal: Who to Trust and What to Believe: Foundations and Methods for Truth-Discovery

Joe Singleton

Given an input consisting of conflicting claims from multiple sources of unknown trustworthiness and reliability, truth-discovery algorithms aim to evaluate which claims should be believed and which sources should be trusted. The evaluations of trust and belief should cohere with one another, so that a claim receives a high belief ranking if it is backed up by trustworthy sources and vice versa.

The design of such algorithms has received an increasing amount of attention in recent years [9, 12, 16, 17], especially with regard to aggregation of information on the web [16] and crowd-sourced data [6, 10]. However the emphasis has been on practical aspects (speed, efficiency etc) rather than theoretical foundations.

In my PhD I will develop a formal mathematical framework for truth-discovery. This will provide a deep understanding of what truth-discovery is in a precise sense, and will facilitate both comparison between truth-discovery and related areas in the literature, and comparison between different truth-discovery algorithms.

Related areas include social choice theory [4, 18], judgement aggregation [7], ranking and recommendation systems [2, 3, 8, 14], argumentation theory [5, 15], and belief revision [13]: at a high level, each of these areas deal with resolving conflicts in data. An interesting question is where truth-discovery fits in this 'big picture'; for example, is truth-discovery a special case of any of these problems, and can results in these areas be applied to truth-discovery? Answering such questions could facilitate development of new truth-discovery techniques using ideas from other areas. By the end of my PhD I aim to have mapped out this landscape.

An approach that has seen great success in the theory of social choice and other areas is the *axiomatic approach*, where desirable properties (called axioms) of voting rules are stated, and voting rules are compared with respect to these properties. K. Arrow's famous impossibility theorem [4] shows that it is impossible for a voting rule to simultaneously satisfy a few certain reasonable

axioms, thus proving a fundamental limitation of voting rules.

Continuing the work started in my final-year undergraduate project, I will develop axioms for truth-discovery, and apply these to real-world algorithms. This will allow algorithms to be evaluated with respect to their theoretical properties as opposed to purely practical ones, and may reveal equivalences between algorithms or uncover fundamental limits of truth-discovery. Another interesting task would be to find a set of sound and complete axioms for a particular algorithm, as has been done in the literature for ranking systems (e.g. in [1] for PageRank [11]).

The project also will have a specific focus on argumentation and its relation to truth-discovery, making use of the school's expertise in this area. In particular, I will investigate how truth-discovery can be represented in an abstract argumentation framework[1], and how techniques in argumentation can be applied to truth-discovery problems (for example, to develop *explainable* truth-discovery algorithms).

Finally, to complement the theoretical work, I will extend the software framework for truth-discovery developed for my final-year project. This will allow me to empirically test whether algorithms satisfy the developed axioms, and will facilitate testing and evaluation of new algorithms.

Prospective supervisors

This proposal falls squarely in line with the research interests of the proposed main supervisor, **Dr Richard Booth**, namely belief change, computational social choice and argumentation theory. Dr Booth has many years of experience working with the axiomatic method, which will be a major component of this project. This proposal has been developed in close consultation with Dr Booth, who is supervisor for my final-year project which this proposal is an extension of. The proposed second supervisor is **Dr Martin Caminada**, who is a leading researcher in argumentation theory and, in particular, dialogue systems, which will be enormously useful for the argumentation parts of the project.

References

[1] Alon Altman and Moshe Tennenholtz. Ranking systems: The pagerank axioms. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, EC '05, pages 1–8, New York, NY, USA, 2005. ACM.

[2] Alon Altman and Moshe Tennenholtz. Axiomatic foundations for ranking systems. *J. Artif. Int. Res.*, 31(1):473–495, March 2008.

[1] A simple approach might include arguments such as 'the value of variable X is x because source s said so'. A counter-argument could be 's is not trustworthy, since they disagree with sources u and v on variable Y'.

[3] Reid Andersen, Christian Borgs, Jennifer Chayes, Uriel Feige, Abraham Flaxman, Adam Kalai, Vahab Mirrokni, and Moshe Tennenholtz. Trust-based recommendation systems: An axiomatic approach. In *Proceedings of the 17th International Conference on World Wide Web*, WWW '08, pages 199–208, New York, NY, USA, 2008. ACM.

[4] Kenneth J. Arrow. Social choice and individual values. *Ethics*, 62(3):220–222, 1952.

[5] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. Abstract argumentation frameworks and their semantics. In Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors, *Handbook of Formal Argumentation*, chapter 4. College Publications, 2018.

[6] Yang Du, Yu-E Sun, He Huang, Liusheng Huang, Hongli Xu, Yu Bao, and Hansong Guo. Bayesian co-clustering truth discovery for mobile crowd sensing systems. *IEEE Transactions on Industrial Informatics*, 2019.

[7] Ulle Endriss. Judgment aggregation. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17. Cambridge University Press, New York, NY, USA, 1st edition, 2016.

[8] Omer Lev and Moshe Tennenholtz. Group recommendations: Axioms, impossibilities, and random walks. In *TARK*, 2017.

[9] Yaliang Li, Qi Li, Jing Gao, Lu Su, Bo Zhao, Wei Fan, and Jiawei Han. Conflicts to harmony: A framework for resolving conflicts in heterogeneous data by truth discovery. *IEEE Transactions on Knowledge and Data Engineering*, 28(8):1986–1999, Aug 2016.

[10] Fenglong Ma, Yaliang Li, Qi Li, Minghui Qiu, Jing Gao, Shi Zhi, Lu Su, Bo Zhao, Heng Ji, and Jiawei Han. Faitcrowd: Fine grained truth discovery for crowdsourced data aggregation. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '15, pages 745–754, New York, NY, USA, 2015. ACM.

[11] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical Report 1999-66, Stanford InfoLab, November 1999. Previous number = SIDL-WP-1999-0120.

[12] Jeff Pasternack and Dan Roth. Knowing what to believe (when you already know something). In *Proceedings of the 23rd International Conference on Computational Linguistics*, COLING '10, pages 877–885, Stroudsburg, PA, USA, 2010. Association for Computational Linguistics.

[13] Gärdenfors Peter. *Belief Revision*. Cambridge Univserity Press, 2003.

[14] Gunawardana A. Shani G. Evaluating recommendation systems. In Ricci F., Rokach L., Shapira B., and Kantor P., editors, *Recommender Systems Handbook*. Springer, Boston, MA, 2011.

[15] Yuqing Tang, Kai Cai, Elizabeth Sklar, Peter Mcburney, and Simon Parsons. A system of argumentation for reasoning about trust. 01 2010.

[16] Xiaoxin Yin, Jiawei Han, and Philip S. Yu. Truth discovery with multiple conflicting information providers on the web. *IEEE Transactions on Knowledge and Data Engineering*, 20(6):796–808, June 2008.

[17] Liyan Zhang, Guo-Jun Qi, Dong Zhang, and Jinhui Tang. Latent dirichlet truth discovery: Separating trustworthy and untrustworthy components in data sources. *IEEE Access*, 6:1741–1752, 2018.

[18] William S. Zwicker. Introduction to the theory of voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 2. Cambridge University Press, New York, NY, USA, 1st edition, 2016.