

The wisdom_of_crowds: an efficient, philosophically-validated, social epistemological network profiling toolkit

Colin Klein¹, Marc Cheong², Marinus Ferreira³, Emily Sullivan⁴, and Mark Alfano³

¹ The Australian National University, ACT, Australia
colin.klein@anu.edu.au

² University of Melbourne, VIC, Australia
marc.cheong@unimelb.edu.au

³ Macquarie University, NSW, Australia

⁴ Eindhoven University, Eindhoven, The Netherlands

Abstract. The epistemic position of an agent often depends on their position in a larger network of other agents who provide them with information. In general, agents are better off if they have diverse and independent sources. Sullivan et al. [19] developed a method for quantitatively characterizing the epistemic position of individuals in a network that takes into account both diversity and independence; and presented a proof-of-concept, closed-source implementation on a small graph derived from Twitter data [19]. This paper reports on an open-source reimplementation of their algorithm in Python, optimized to be usable on much larger networks. In addition to the algorithm and package, we also show the ability to scale up our package to large synthetic social network graph profiling, and finally demonstrate its utility in analyzing real-world empirical evidence of ‘echo chambers’ on online social media, as well as evidence of interdisciplinary diversity in an academic communications network.

Keywords: social epistemology, Python, social network analysis, testimonial networks

Note About This Preprint

This is a preprint of the following chapter: Colin Klein & Marc Cheong & Marinus Ferreira & Emily Sullivan & Mark Alfano, “The wisdom_of_crowds: an efficient, philosophically-validated, social epistemological network profiling toolkit”, published in “Complex Networks & Their Applications XI: Proceedings of The Eleventh International Conference on Complex Networks and their Applications: COMPLEX NETWORKS 2022 - Volume 1”, edited by Hocine Cherifi, Rosario Nunzio Mantegna, Luis M. Rocha, Chantal Cherifi, Salvatore Micciche, 2022, Springer reproduced with permission of Springer.

The final authenticated version is available online at: https://link.springer.com/chapter/10.1007/978-3-031-21127-0_6.

1 Introduction

Most of what we know we know because we learned about it from other people. *Social epistemology* is the subfield of philosophy that studies how knowledge and justification depend on the testimony of others transmitted through social networks [7]. A focus on networks has been influential because it allows philosophers to connect their concerns to the substantial body of empirical and simulation work on real-world networks and their graph-theoretic properties.

Sullivan et al. [19] presented a method for quantitatively characterizing the epistemic position of individuals in a network. Broadly speaking, individuals are in a better epistemic position if they are receiving information from *diverse* and *independent* sources, with the more diversity and independence the better. This method was based on, amongst others, the *Wisdom of Crowds* hypothesis, that the aggregated judgements of many individuals can systematically be more accurate than the judgements of those individuals taken singly [20]. They then operationalized these two concepts in a way that allowed them to provide an interesting profile of a small 185-member Twitter community [19]. That work relied on a bespoke, closed-source codebase. As it was built as a proof-of-concept, it was also not optimized in ways that naturally scaled to larger networks. This made it difficult to apply the technique to other datasets, such as networks from other social media sites, or networks created from artificial social simulation algorithms, e.g. Laputa [16]: all of which are of interest to both philosophers and computer scientists alike. That work’s codebase had an emphasis on the generation of Java-based visualizations—using a combination of several platforms and toolchains— which does not lend itself to convenient large-scale network analysis, due to performance limitations.

To make this tool more widely available to researchers, we therefore present `wisdom_of_crowds`, a complete ground-up re-implementation in Python of the core Sullivan et al. [19] concepts. The code is optimized to deal with larger networks. It also includes some standardized helper functions to allow for coordinating results between research groups and data scientists. We have made the code for this package open source, under the *GNU General Public License 3.0*⁵, on GitHub (<https://github.com/cvklein/wisdom-of-crowds/>). It has also been accepted to the Python Package Index (PyPI, at <http://pypi.org/project/wisdom-of-crowds>), and is available for any user to install via `pip`⁶.

As much as possible, we have relied on open-source Python packages such as `networkx` [17] and `matplotlib` [9] as they have been rigorously tested and are freely available for auditing and peer review. For good practices in verification and validation, the `pytest` [10] library is used to provide a unit-testing framework. Having created this new implementation, we sought out to deploy it in our investigation of contemporary social network data, to provide a data-driven perspective to complement existing conceptual and theoretical work.

This paper thus aims to present two key findings:

⁵ Full license terms can be found at <https://www.gnu.org/licenses/gpl-3.0.en.html>.

⁶ It can be installed by the Python community via: `pip install wisdom_of_crowds`

- the core concepts of the `wisdom_of_crowds` package, including compatibility with (and optimization for) contemporary network-related datasets and packages in Python. This includes improvements to the base [19] algorithm, by clearly defining the bounds (and justifications for) parameters used, as well as suggested extensions including the h -measure derived from [8].
- application of `wisdom_of_crowds` on simulated large networks; to investigate its feasibility/performance
- application of `wisdom_of_crowds` on actual real-world networks—Twitter discourse about the Black Lives Matter (BLM) movement between January and July 2020 centering around the May 25th 2020 murder of George Floyd [2]; and the *email-Eu-core* network on communications patterns in “a large European research institution” [11,23,12]—to corroborate theoretical network epistemological findings on modern-day social networks centered on current phenomena.

2 Background and Methods

2.1 Core concepts

Epistemology is a branch of philosophy which, simply put, “is concerned with how people should go about the business of trying to determine what is true” [18]. As per our Introduction (Section 1), *social epistemology* concerns the testimony of others embedded in social contexts [7]; in contrast with ‘individual’ epistemology which concerns how an individual conducts reasoning, abstracted away from their “social environment” [18].

In recent years, social epistemologists have moved away from considering dyadic relationships between individuals to consider the ways in which social epistemic *networks* shape the information we receive [15,4]. Consider an epistemic network G where nodes are epistemic agents and edges represent the relationship of receiving information via testimony. ‘*Testimony*’ is used broadly in social epistemology for any way in which one source delivers information to another, and includes speech, writing, and other forms of media. All things being equal, a node is better off receiving information from more and more diverse nodes. However, testimony is often transmitted in chains, and this transmission need carry only the content of the information, not (meta-)information about the original source or the intermediate links. This complicates the position of any individual who is trying to learn from multiple sources. For example, a piece of gossip heard from two people seems more reliable than from one, but that reliability is undermined if both heard it from the same person [3].

2.2 Sullivan et al. (2020)’s operationalizations

Defining the m, k observer: Following [19], we say that a node n is an m, k -observer just in case it receives information from a set of at least k different nodes which are pairwise at least m steps away from one another, when considered on

the subgraph of G that does not contain n . If G is directed, then candidate sources must be at least m steps away in both directions. The removal of n from consideration in the case of distances is necessary for directed graphs, as otherwise all sources to n are trivially at most 2 steps apart; we carry over that requirement to undirected graphs as well.

In this work, we assume $1 \leq m \leq 5$, as most real life networks have length 6 paths between most arbitrarily chosen nodes [13]. We bound $2 \leq k \leq 5$, because a node with a single source is in a very poor epistemic position with respect to diversity of input. Note that it is a consequence of the definition that if n is an m, k -observer, it is also both an $m - 1, k$ -observer and an $m, k - 1$ -observer (assuming $m - 1$ and $k - 1$, respectively, are permissible values).

Given this definition, the core concepts in [19] are defined as follows.

$S(n)$, independence of sources. $S(n)$ gives a measure of the independence of sources to node n . Consider the set s of possible pairs (m, k) for which n is an m, k -observer. Then define

$$S(n) = \begin{cases} 0 & \text{if } s = \emptyset \\ \max\{mk : (m, k) \in s\} & \text{otherwise} \end{cases} \quad (1)$$

In other words, $S(n)$ is just the largest mk such that the n is an m, k -observer. If S has 0 or 1 nodes as sources, they are considered as being in an epistemically bad position, and so $S(n) = 0$. Note that given this definition, possible S values do not increase smoothly. Given the bounds set out above, $S(n) \in \{0, 2, 3, 4, 5, 6, 8, 9, 10, 12, 15, 16, 20, 25\}$.

$D(n)$, diversity of sources. $D(n)$ measures the diversity of the sources that contribute information to n . Let each node i be associated with a set a_i of epistemically-relevant attributes. These might be group affiliations, topics of interest, scientific approaches, political leanings, and so on. Let s be the set of n 's sources. Then define

$$D(n) = |\bigcup\{a_i : i \in s\}| \quad (2)$$

That is, D gives the number of distinct *types* of information that feed into n .

$\pi(n)$, epistemic position. Finally, as the epistemic position of a node is a function of both the diversity and independence of sources, we define

$$\pi(n) = S(n)D(n) \quad (3)$$

2.3 Caveats and Considerations

There are a few notes to make about the implications of the Sullivan et al. [19] core concepts in real social networks.

Regarding m, k -observers, while higher rankings are better and all the nodes with a specific value of m and k are members of an equivalence class, the framework does not posit which of two m, k -observers is better positioned if one has a higher m value and the other a higher k value—for instance, whether 2,3-observers are better- or worse-placed than 3,2-observers. The framework thus does not provide a total order but instead provides a collection of *partial* orders.

Intuitively, for instance, a 1,2-observer is worse placed than a 2,3-observer despite having neither an m or a k value in common, because a 2,3-observer is better placed than a 2,2-observer when considering k values, which in turn is better placed than a 1,2-observer when comparing m values. Thus, a node’s S value does provide a way to more easily compare the epistemic position of nodes, by combining their m and k values into a single value. There are concerns with it, though.

The choice of multiplying m and k values together to come to a single measure is ultimately arbitrary. The S value contains less information than the m and k values does, because it fails to preserve the difference between being a 2,3-observer and being a 3,2-observer (the former has fewer independent sources, but these have a greater degree of independence from each other than in the latter’s). Simultaneously, it posits that a node which is a 3,3-observer is determinately better placed than either a 4,2-observer or a 2,4-observer, but worse than a 5,2-observer or a 2,5-observer. It is unclear why we should believe that this is true in general.

In short, while a node’s S value is one measure of the independence of its sources, it is unclear why we should use this measure rather than another. This was the impetus for us introducing a further measure, the `h_measure(n)` [8] which returns the highest h such that n is a h, h -observer, as discussed below.

2.4 Our Re-Implementation

The core of the `wisdom_of_crowds` is a class `Crowd`⁷. `Crowd` is initialized with a `NetworkX` graph (encapsulating the social network’s edges and nodes), and provides various functions to calculate the metrics defined above.

Calculating whether a node is an m, k -observer combines multiple shortest-path problems with clique-finding problems. Naïve approaches to the latter have complexity $O(n^k)$ [22]. Given that we are considering unweighted paths, the shortest-path problem has a reasonably efficient linear-time solution, but the requirement to remove n from the distance calculations also means that network shortest paths cannot simply be calculated at the outset. In the worst case scenario, they must be recalculated for each pair of sources for each node.

Hence this is a computationally difficult problem to brute-force. In practice, efficient caching and testing of seen paths plus greedy k -clique algorithms means that worst-case performance can often be avoided for realistic networks:

⁷ Due to space constraints, we only discuss our substantial contributions within this paper. The full documentation for all implemented methods is available at https://github.com/cvklein/wisdom-of-crowds/blob/main/docs/wisdom_of_crowds.py.md

we *memoize* (i.e., cache) intermediate path values, trading-off space for a reduced processing time. As the envisioned use case for large graphs is for one-shot batch processing, our code allows for easy multithreading or multiprocessing (e.g. using *multiprocessing* or *concurrent.futures*) allowing it to attain substantial performance gains, in conjunction with the memoization mechanism.

The m, k -observer functionality is the basis for calculating D , S , and π . D is calculated on node attributes, and users can specify the appropriate key for the attribute. If a single attribute is supplied, D is calculated using the singleton set containing that attribute.

In addition to the standard measures, we also introduce an improvement: the `h_measure()` of a node n as the smallest x such that n is an x, x -observer; comparable to the standard definition of Hirsch’s h -index in citation practices [8]. As suggested by [19], being a 3,3-observer is the minimal secure epistemic position, and the use of a single non-multiplicative standard may be useful for some cases.

Finally, the package includes two helper functions to allow for comparable reporting and display across different users. `iteratively_prune_graph()` takes a *NetworkX* graph, removes small-degree nodes, small-weight edges, and takes the largest connected component in what remains, iterating this process until the graph is stable. The thresholds are parameterized; the default is for $indegree + outdegree \leq 1$ and no edge culling, as per [19]. In the spirit of scientific reproducibility, `make_sullivanplot()` makes a summary figure of a whole network in the style of [19, refer to their Figure 7]. It can produce standalone plots or return a subplot in a specified `matplotlib` axis.

3 Experimental Results

3.1 Efficiency Tests

Firstly, to benchmark the ability of `wisdom_of_crowds` on different magnitudes of nodes, typical of modern-day datasets, we sought to batch calculate S for different magnitudes of node sizes ($|N|$), given various probabilities of edge connections.

Figure 1 shows the efficiency of batch calculating S for each node of a **Crowd** on random graphs with varying parameters for probability of edge connection. Random directed graphs were generated using the `networkx` generator `fast_gnp_random_graph()` with the parameters indicated in the figure. Timing was done using the python `timeit` package over a single iteration.

As the log-linear plot in Figure 1 shows, there is a roughly exponential relationship between the number of nodes and runtime, with the exponent a function of the number of edges. This suggests that the efficiency of our code approaches what would be expected given the fundamental complexity of the clique-finding problem. Note that the exponential growth means that the boundary between computationally tractable and intractable graphs can be relatively tight. Judicious pruning often makes a difference.




Fig. 1. Timing curve for random graphs by number of nodes. Different markers represent different connection probabilities for nodes.

Further, note that this was achieved in standard Python operating conditions, i.e. in the absence of any multiprocessing/multithreading support (see also Section 2.4). Performance gains will be attained for batch calculations on high-CPU (16-or-more CPU cores) systems.

3.2 Application on Social Network Data: BLM on Twitter

Secondly, as part of our investigation of real-world phenomena, in line with [19], we replicate their findings on a real-world Twitter retweet network to examine the information-sharing dynamics during the Black Lives Matter movement. An earlier version of this dataset was used by [19] who were able to examine a network of 185 nodes. For further information on the dataset⁸ used, see [2]. For brevity, the key details are summarised here.

A *retweet network* [19,2] was generated: i.e., a weighted directed network where nodes are authors and the weight of an edge from node u to node v represents the number of times that user v retweeted user u . We took the largest connected component of this graph as the starting point for cluster-analysis [21,6]. Following [2], We found first-level clusters using Modularity Vertex Partitioning, preserving clusters with more than 10% of the original nodes. This gave 4 clusters, covering 83% of the graph. Next, we manually inspected the 100 most-influential nodes within each group, characterizing the communities

⁸ We queried the Twitter Streaming API with a series of Black Lives Matter (BLM)-related keywords, hashtags, and short expressions in a window between January and July 2020. The dataset comprised ~ 4.6 M original tweets between January 13th and July 18th and ~ 94.5 M retweets from January 18th to July 23rd; these tweets were produced by ~ 2.0 M distinct authors.

as Activists, Center-Left Democrats, Republicans, and a set of “Boosters” who mainly amplified the content of the first two groups. Topic models were fit using *scikit-learn*’s non-negative matrix factorization (NMF), fit on a *tf-idf* representation of the Twitter text (post-sanitization) with `min_df=0.05`. Finally, we used `iteratively_prune_graph()` with a node and weight threshold of 3, which resulted in a tractable subgraph with $|N| = 16249$ nodes and $|E| = 145246$ edges. This subgraph had very little representation from the ‘booster’ group, so they were omitted from further analysis. By comparison with [19], our analysis is an increase in the number of nodes by a magnitude of $\sim 100\times$. Batch processing took about 6.25 hours on a 2017 desktop iMac.

Figure 2 plots S , D , and π for the current experiment. For the left side of 2, D was calculated using the cluster identity of the node. For the right side, D was calculated using the `argmax` of the fitted and normalized W matrix for the topic model. This gives the topic that is most distinctive of each user’s tweets. We examine both the network as a whole and three identified subclusters in the graph. The left half of Figure 2 shows S , D , and π for the network, where D is calculated via the subcluster identity of sources. The right half of Figure 2 recalculates D and π based on a 9-topic NMF topic model of aggregated tweets (compared with the 3-topic model of [19]).

The results illustrate the utility of profiling networks using our toolkit. On the left, one can see that Republicans appear to be in the worst epistemic position in terms of the other subgroups with which they interact: they have a generally low D , suggesting that they tend to listen mostly to in-group members. However, they have a relatively high S and therefore a π comparable to other subgroups. Compare this to the topic-wise graph, in which Republicans have a relatively high diversity for *topics*, one at least as good as other groups. The activist group shows something of the inverse pattern. That is, they show a more varied range of S and D values for group-group interactions, but a comparatively lighter graph with fewer topics for the broad span.

These results might suggest that both groups are part of ‘echo chambers’ [1,14], but in different ways: the right tends to be a monoculture socially but a polyculture topically, with a converse pattern on the left. Finally, we note that all subgroups, in both domains of measurement, tend to have an $S < 10$ for more than half the population. This replicates the observations of [19], in which most participants end up in a comparatively poor epistemic position. However, most groups tend to contain at least a small sub-population which is well-connected and often with a relatively high D . We note that this is especially the case with our ‘Booster’ group, a small subset who seemed to be mainly concerned with amplifying and relaying the content of other groups.

3.3 Application on Communication Network Data: *email-Eu-core*

Finally, we apply our `wisdom_of_crowds` analysis onto the *email-Eu-core* network, an existing dataset curated by [12,23] on the Stanford Large Network Dataset Collection [11]. The rationale behind this is to apply our social epistemo-

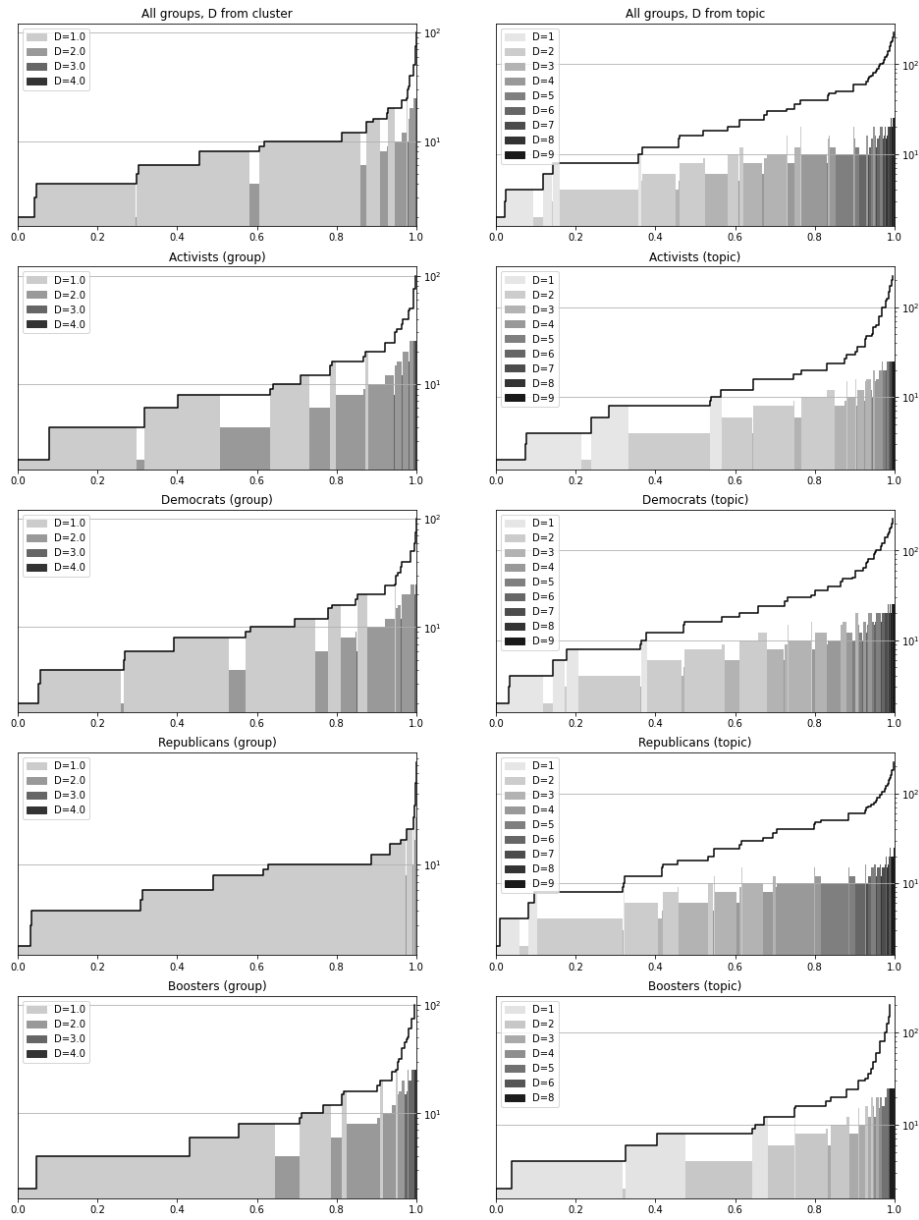


Fig. 2. Profile plots for entire network and subgroups looking at clusters (left) and topics (right). X axis is proportion of total, Y axis shows both S (height of bars) and π (black line), plotted on a log scale.

logical lens to analyze an existing network which is hitherto (to our knowledge) never been examined using the tools we have at hand.

Briefly, *email-Eu-core* consists of “email data from a large European research institution” [11], represented as a digraph where an edge (u, v) exists “if [researcher] u sent [researcher] v at least one email”. The beauty of *email-Eu-core* lies in the fact that it only considers internal communications in the institution, ignoring any noise resulting from possible non-academic communication originating from/in reply to outside actors; and that the ground-truth membership of each node has already been established, i.e., as belonging to any “one of 42 departments at the research institute” [11].

Compared to more modern social networks, *email-Eu-core* has a comparatively small $|N| = 1005$ and $|E| = 25571$. However, again, we note that this is already $\sim 10\times$ the magnitude as compared to the hitherto state-of-the-art [19]. The total running time on such a dataset is comparatively trivial (less than 5 minutes). As we do not have e-mail text associated with the original nodes in this dataset, we sought only to analyze the ‘overall epistemic picture’ of the network, as a whole. In the same vein as Section 3.2, Figure 3 illustrates the profile plot for *email-Eu-core*.

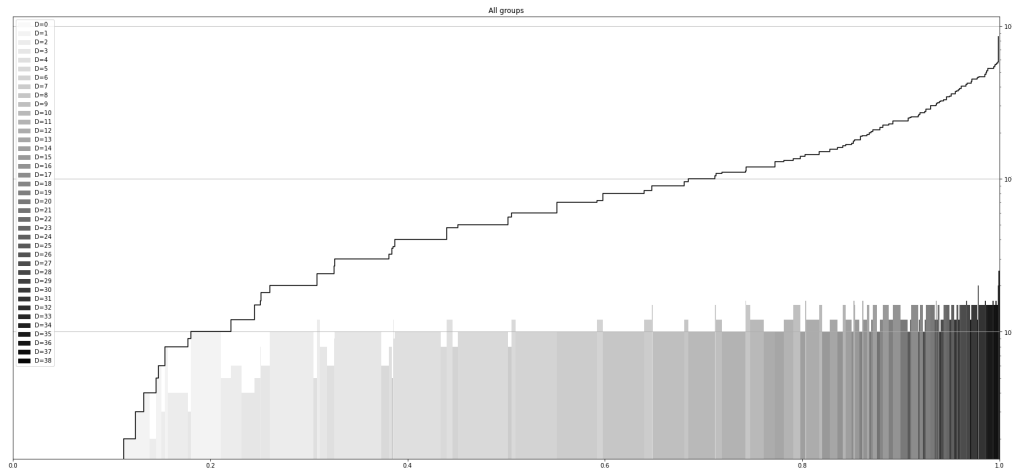


Fig. 3. Profile plots for the entire *email-Eu-core* network [12,23,11]. The X axis is proportion of total, Y axis shows both S (height of bars) and π (black line), plotted on a log scale.

Again, the utility of our approach is evident. Given the results, we observe that the distribution of S is fairly consistent for all the nodes (researchers) in the research institution’s academic network. However, the progressively darker bars illustrate that researchers who have connections with a more diverse amount of departments (thereby maximizing D) can vastly optimize their epistemic position. Roughly, about $\sim 40\%$ of researchers – i.e., the right-most data points in Figure 3 – have a π of about 100 or more, despite having roughly the *same* S values.

To our knowledge, this is the first time empirical social epistemological analyses in the spirit of [19] have been conducted on such email networks.

4 Discussion

Our results show that it is possible to replicate the methodology used by [19] to larger networks, and that insights about the relative epistemic positions of different communities within a network can be drawn from plotting these parameters. As our package and its dependencies are all open source, this makes it possible for researchers in a range of fields (including philosophy, psychology, sociology, anthropology, communications, and network science) both to conduct new research and to reanalyze networks that they have previously studied.

We anticipate that future research will expand the types of social networks under study. Other sources from online social media such as Facebook, Reddit, and YouTube all seem to be viable candidates for study. Considering offline epistemic networks would be especially valuable, such as those found in the landmark Bernard-Kilworth-Sailer (*BKS*) analyses of social networks [5], as their structure may be interestingly different from the structures found online. `wisdom_of_crowds` also allows us to conduct research on epistemic network simulations, created with tools such as Laputa [16], which can quickly scale up to thousands of nodes. We expect that studies of friend networks, organizational networks in industry and the military, networks of sources used by journalists, criminal cartel networks, and academic citation networks would prove valuable.

Moving beyond that, it would be interesting to study networks with more than one type of testimonial edge (e.g., public communications versus private ones). One intriguing hypothesis is that these may differ in structure even if they contain the same nodes, and that individuals who are central in public networks but peripheral in private networks (or vice versa) would tend to play unique roles in the social epistemology of those networks. For instance, someone who is privately in communication with a very large number of others but not publicly visible is in a position to exert influence because the others may assume that they have a much better epistemic position than they actually do.

The exploratory profiling made possible by our tool reveals patterns of epistemic isolation and interaction across real-world networks, and suggests possibilities for more specific analyses. By providing it to the community at large, we hope to facilitate further modeling of epistemic networks across a variety of domains.

Funding

Work on this paper was supported by ARC Grant DP190101507 (to C.K. and M.A.) and by Templeton Grant 61378 (to M.A.).

This research was supported by use of the Nectar Research Cloud and by Melbourne Research Cloud (University of Melbourne) (to M.C.). The Nectar

Research Cloud is a collaborative Australian research platform supported by the NCRIS-funded Australian Research Data Commons (ARDC).

References

1. Alfano, M., Carter, J. A., and Cheong, M. (2018). Technological seduction and self-radicalization. *Journal of the American Philosophical Association*, 4(3):298–322.
2. Alfano, M., Reimann, R., Quintana, I. O., Chan, A., Cheong, M., and Klein, C. (2022). The affiliative use of emoji and hashtags in the Black Lives Matter movement on Twitter. *Social Science Computer Review*, Accepted for publication (forthcoming).
3. Alfano, M. and Robinson, B. (2017). Gossip as a burdened virtue. *Ethical Theory and Moral Practice*, 20(3):473–487.
4. Alfano, M. and Sullivan, E. (2020). Humility in social networks. In *The Routledge Handbook of Philosophy of Humility*, pages 484–494. Routledge.
5. Bernard, H. R., Killworth, P. D., and Sailer, L. (1979). Informant accuracy in social network data IV: a comparison of clique-level structure in behavioral and cognitive network data. *Social Networks*, 2(3):191–218.
6. Csardi, G. and Nepusz, T. (2005). The igraph software package for complex network research. *InterJournal, Complex Systems*:1695.
7. Goldman, A. I. (1999). *Knowledge in a social world*. Oxford University Press.
8. Hirsch, J. E. (2005). An index to quantify an individual’s scientific research output. *Proceedings of the National Academy of Sciences*, 102(46):16569–16572.
9. Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3):90–95.
10. Krekel, H., Oliveira, B., Pfannschmidt, R., Bruynooghe, F., Laugher, B., and Bruhin, F. (2004). pytest 6.2.5.
11. Leskovec, J. (n.d.). Snap: Network datasets: email-eu-core network. <https://snap.stanford.edu/data/email-Eu-core.html>.
12. Leskovec, J., Kleinberg, J., and Faloutsos, C. (2007). Graph evolution: Densification and shrinking diameters. *ACM Trans. Knowl. Discov. Data*, 1(1):2–es.
13. Milgram, S. (1967). The small world problem. *Psychology today*, 2(1):60–67.
14. Nguyen, C. T. (2020). Echo chambers and epistemic bubbles. *Episteme*, 17(2):141–161.
15. O’Connor, C. and Weatherall, J. O. (2019). *The misinformation age*. Yale University Press.
16. Olsson, E. J. (2011). A simulation approach to veritistic social epistemology. *Episteme*, 8(2):127–143.
17. Schult, D. A. (2008). Exploring network structure, dynamics, and function using networkx. In *In Proceedings of the 7th Python in Science Conference (SciPy)*. Citeseer.
18. Steup, M. and Neta, R. (2020). Epistemology. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy (Fall 2020 Edition)*. <https://plato.stanford.edu/archives/fall2020/entries/epistemology/>.
19. Sullivan, E., Sondag, M., Rutter, I., Meulemans, W., Cunningham, S., Speckmann, B., and Alfano, M. (2020). Vulnerability in social epistemic networks. *International Journal of Philosophical Studies*, 28(5):731–753.
20. Surowiecki, J. (2005). *The Wisdom of Crowds*. Abacus, London.
21. Traag, V., Waltman, L., and van Eck, N. J. (2019). From Louvain to Leiden: Guaranteeing well-connected communities. *Scientific Reports*, 9:5233.

22. Vassilevska, V. (2009). Efficient algorithms for clique problems. *Information Processing Letters*, 109(4):254–257.
23. Yin, H., Benson, A. R., Leskovec, J., and Gleich, D. F. (2017). Local higher-order graph clustering. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, page 555–564, New York, NY, USA. Association for Computing Machinery.