



DOI:10.1145/3588431

BY DAMIANO SPINA, MARK SANDERSON,
DANIEL ANGUS, GIANLUCA DEMARTINI, DANA MCKAY,
LAUREN L. SALING, AND RYEN W. WHITE

Human-AI Cooperation to Tackle Misinformation and Polarization

A DOMINANT NARRATIVE of the past decade is that algorithms contribute to a misinformed and segregated society. Perhaps paradoxically, algorithms are often sought as solutions to such problems. We describe a significant emerging trend away from this techno-solutionist approach that seeks to create and understand a new paradigm: a productive interplay between algorithms and people. Two relevant test cases are being explored in our region: The first addresses a new framework to tackle misinformation by assisting fact-checkers with computational methods, and the second seeks new models to understand how search engines deliver personalized search results when little or no algorithmic personalization exists.

In late 2020 and early 2021, the Australian Communication and Media Authority conducted a


study to analyze the state of misinformation in Australia. The findings, reported to the Australian Government in June 2021, showed that four out of five Australian adults had been exposed to misinformation about COVID-19. They also found that online misinformation, such as the propagation of anti-vaccine narratives within the Australian community, had a direct negative impact on the trust that people place in democratic institutions and public health agencies. These narratives often originate overseas but quickly spread through local communities. The fact-checking organizations that have traditionally verified statements made by public figures or politicians in public and mainstream media now must also monitor and debunk dramatically faster-spreading claims on social media platforms. Narratives containing misinformation are having a direct and negative impact on how people consume information: They may influence the content we engage with and the search terms we enter.¹⁰ Given that an informed citizenry is a cornerstone of democracy, public decision making is at risk.

The significance of the problem was also recognized in the International Cyber and Critical Technology Engagement Strategy released by the Australian Government, which identifies digital misinformation as a clear risk to the security and safety of Australia, the Indo-Pacific region, and beyond. Countries across East Asia and Oceania introduced legislation that specifically targets so-called ‘fake news’ and they created voluntary codes of practice developed in partnership with the technology industry.

Despite such efforts, as of December 2022, out of the 122 currently verified signatories of the Poynter’s International Fact-Checking Network (IFCN), only eight



Computational methods to assist in fact-checking.



are in the East Asia and Oceania region: Australian Associated Press (AAP) and RMIT FactLab in Australia; Cek Fakta Liputan 6, MAFINDO, Tempo.co, and Tirto ID in Indonesia; and Rappler and Verafiles Incorporated in The Philippines.^a

Some platforms have turned to fact-checkers to help identify problematic content. However, the deluge of misinformation means that the checkers cannot handle the large number of claims that need to be assessed. Algorithmic assistance may therefore be beneficial to identifying instances of misinformation.

Computational Methods to Tackle Misinformation

In the last few years, a trend has emerged of computing professionals leveraging hybrid-intelligence (human and artificial) methods to remove misinformation from online platforms.⁶ The problem is more complex than identifying and removing misinformation; such content needs to be comprehensively managed throughout the stages of its lifecycle: from creation to propagation and consumption. The computer science community has already developed technologies that can help at each stage (see the

accompanying figure).

A range of methods—including social network analysis, natural language processing (NLP), information retrieval (IR), knowledge graphs, machine learning (ML) and artificial intelligence (AI), foundation models (for example, neural transformers such as BERT and GPT) and deep learning, data visualization, explanations of machine-learned model output, and advances in human-computer interaction (for example, new user experiences and interaction capabilities)—can assist but not replace humans during misinformation management processes. In all these stages, a close collaboration between experts, systems, and non-experts such as crowd workers is crucial to scaling up while maintaining the quality, agency, and accountability of the process.⁶

Human-in-the-loop fact-checking in the East Asia and Oceania region is in its infancy. While non-governmental organizations such as FirstDraft News (now the Information Futures Lab) are active in the region, more research is needed to better understand how hybrid-intelligence methods can be effectively embedded into misinformation management processes without taking agency away from experts.⁹ In addition

to debunking misinformation, computational methods can assist in *prebunking* processes by developing effective ways to educate people about misinformation, thus enhancing digital literacy. This will provide them with the skills to identify and question unverified information online. The low number of fact-checking organizations in East Asia and Oceania makes the support of computational methods more pressing in this region. Ensuring that this algorithmic assistance is useful, though, will also require region-specific attention.

What is ‘fact’ internationally is often counterfactual in East Asia and Oceania, so merely importing international information will not be effective. Examples of ‘facts’ that do not hold true in this region include seasonal issues; for much of the region summer is December–February. While having Christmas in summer is slightly countercultural, gatherings for these holidays did result in a major southern summer COVID-19 spike in 2021. Conversely, knowing that the influenza (and COVID-19) seasons occur in June and July in this part of the world is a key part of good public health advice. The region also predominantly sits close to the equator, so public health advice about sun exposure must be tailored. Most countries

^a <https://ifcncodeofprinciples.poynter.org/signatories>


in the region drive on the left-hand side of the road, affecting road safety advice. There are also cultural differences within the region: The weekend can fall on different days, or many countries are predominantly Muslim, meaning Eid, not Christmas, is celebrated. Understanding the importance of COVID-19 vaccines being Halal was key to public health messaging in Melbourne, Australia. Democratic conventions are also unique: Australia has one of the highest rates of democratic participation in the world, at more than 90%, a direct result of compulsory voting. Nearby New Zealand, though, also had strong democratic participation, more than 80% without compulsory voting in the most recent election. Given these regional and local differences from global norms, local fact checking is key to ensuring an informed populace. Further, prebunking strategies must sit alongside existing educational and cultural norms. This presents the two-pronged problem of scarce local expertise and the need for localized resources, to build and evaluate algorithmic tools and human-in-the-loop solutions for the region.

The Filter Bubble Myth


While there is evidence that polarization in society dramatically escalated with the introduction of broadband Internet, the cause is not well understood. Filter bubbles, formed by algorithms delivering personalized content that reinforces a particular worldview, have become an incredibly popular explanation. The conceptualization of the filter bubble was coined in a book of the

same name by Eli Pariser.⁷ However, the filter bubble concept may distract from the deeper epistemic causes of polarization. Pariser's book, cited thousands of times, makes the case but provides limited evidence of bubbles being formed by search engines. Empirical studies indicate a lack of such bubbles, going further to suggest that search platforms increase exposure to contrary viewpoints. Cross-disciplinary teams of computer scientists, media specialists, information scientists, industry researchers, and psychologists are working together on the issue of search personalization through novel experimentation, which has better revealed the role that search engines play in polarization.

The Australian Search Experience,³ a project carried out by the Australian Research Council Centre of Excellence for Automated Decision-Making and Society (ADM+S), is a data-donation study where more than 1,000 people across Australia were recruited to examine whether search engines returned different kinds of results across the cohort. Participants installed a Web browser plug-in that issued periodic queries to well-known search engines using the participants' accounts. Search results were scraped by the plug-in and returned to researchers for examination. The queries were drawn from a pre-defined list of common searches on a range of topics spanning political, controversial, and everyday categories. While the research is ongoing, initial findings indicate that, although search results were found to be contextualized to specific




Narratives containing misinformation are having a direct and negative impact on how people consume information.




Sample of query variants crowdworkers generated, drawn from the UQV test collection.²

Description of Information Need	Example Query Variants
A group of local farmers has been protesting outside your energy utility's offices, complaining about a plan to build a wind farm on hills near their properties. Their placards say that the disadvantages (cons) outweighed the advantages (pros). Until now, you had always assumed that wind power was a good thing; now you are not so sure and decide to find out more.	▶ Advantages and disadvantages of wind power
	▶ Cons of wind power
	▶ Is wind power good?
	▶ Negative effects of using wind power
	▶ Pros and cons of wind power
	▶ Information about wind power
	▶ Engineering principles of wind turbines



A trend has emerged of computing professionals leveraging hybrid-intelligence methods to remove misinformation from online platforms.



geographic locales, algorithmic personalization in search engines may be less extensive than was suggested by previous filter-bubble research. This leads to the question: If search is largely homogeneous, where is information polarization coming from?

One possible answer lies in work of our region's IR community examining the impact of query variation in search. While search-engine users have been studied for decades, recent experiments where a large number of people are asked what query they would use when seeking to satisfy a common information need have found an astonishing range of distinct queries.² To illustrate, the accompanying table lists a sample drawn from more than 50 query variants found when 100 crowdworkers were asked how they would search for information about wind power. Such extensive variations were recorded across a diverse set of 100 topics. The results of the experiment are packaged in a test collection that captures this user query variability (UQV).

Query variations were found to have a significant impact on search-engine performance. Wide variations in the queries submitted to commercial search engines were identified,¹ and detailed statistical analysis found that variations in queries had a substantially larger effect on search results than any change in the workings of a search algorithm.⁵ The Australian Associated Press recently debunked a social media post with a false claim about the lifespan of wind farm generators.^b One can see in the sample queries shown in the table that different queries appear to reflect different attitudes on the topic. It is natural to wonder whether misinformation influences the way people choose the keywords that they type into a search engine.

Opportunities to Address Polarization in Search Engines

This collection of findings suggests that polarization in search is being

driven not by algorithms but by searchers.¹ The research trend highlights a critical oversight in search-engine algorithm design: understanding how search algorithms react to and potentially alleviate this user variation. The research challenges of such work include:

- ▶ Understanding the reasons for the variation people show when searching—for example, demographics, search habits, domain knowledge, cognitive biases, and how people are prompted to search.
- ▶ Exploring if and how people are influenced by others to search in particular ways.
- ▶ Determining how search algorithms can be adapted to better handle the variation.

Initial results suggest that the way people construct queries is informed by established searching habits, although other factors, such as existing knowledge, biases, prompts, and so on, most likely also contribute. Questions about how people are influenced to search must be examined. Here, misinformation seems to play a crucial role, and collaborations with fact-checking organizations in the region^{4,8} are helping to better understand how people formulate their queries when they encounter misinformation and interventions—for example, verified content produced by fact-checkers. When examining search algorithms, the roles and responsibilities of search engines will be questioned. Most would agree that search engines should return the most reliable content, but should they intervene in trying to change user views arising from polarizing queries? Answers to such questions must be approached in a nuanced way because confronting people with views too distant to their own could alienate them. Other key research issues include investigating whether queries resulting from a disinformation campaign can be reliably detected, whether search engines could and should detect the sole pursuit of confirmatory information and evidence of confirmation biases,

^b <http://bit.ly/turbineFactCheck>

and whether search results can be tailored to support reflection and understanding of those with beliefs different to one's own.


The Road Ahead

Detailing these two case studies shows that a richer engagement between humans and machines ensures more effective outcomes in the management of information and a better understanding of how online information is sought. The work described here represents an important emerging trend in our region, but the impact is felt far beyond this geographic area. This work also presents several grand challenges in deploying these assistive technologies at a massive scale and realizing human-AI cooperation in practice.^c

Computing professionals must continue collaborating with other disciplines to make and integrate advances in critical areas, such as fairness, accountability, transparency, explainability, and the safety of human-AI cooperation. Misinformation and its exposure will only grow in the coming years, as will the adversarial uses of computational methods to generate and spread disinformation narratives. Polarization will persist as long as we fail to understand the causes of query variation in search-engine engagement and fail to develop more robust search algorithms capable of handling that variation. As a community, we must meet these challenges head on. Understanding and supporting the interplay of humans and algorithmic systems will ultimately lead to better outcomes for all.

Acknowledgments. The authors thank Axel Bruns, Luke Gallagher, Timothy Graham, James Meese, Stefano Mizzaro, Quoc Viet Hung Nguyen, Abdul Karim Obeid, and Falk Scholer for their contributions and feedback toward this work. This research is partially supported by the Australian Research Council

^c Also highlighted in the ACM Technology Policy Council's Statement on Principles for Responsible Algorithmic Systems; <https://bit.ly/jointAIstatement>

(DE200100064, CE200100005, IC200100022). The authors acknowledge the Traditional Custodians of Country throughout Australia and their connections to land, sea, and community. We pay our respect to their Ancestors and Elders past and present and extend that respect to all Aboriginal and Torres Strait Islander peoples today. 

References

1. Alaofi, M. et al. Where do queries come from? In *Proceedings of the 45th Intern. ACM SIGIR Conf. Research and Development in Information Retrieval (2022)*, 2850–2862; <https://doi.org/10.1145/3477495.3531711>.
2. Bailey, P., Moffat, A., Scholer, F., and Thomas, P. UQV100: A test collection with query variability. In *Proceedings of the 39th Intern. ACM SIGIR Conf. Research and Development in Info Retrieval (2016)*, 725–728; <https://doi.org/10.1145/2911451.2914671>.
3. Bruns, A. *Australian Search Experience Project: Background Paper*. Technical Report. ARC Centre of Excellence for Automated Decision-Making and Society (2022); <https://doi.org/10.25916/k7py-t320>.
4. Cerone, A. et al. Watch 'n' Check: Towards a social media monitoring tool to assist fact-checking experts. In *Proceedings of the 2020 IEEE 7th Intern. Conf. Data Science and Advanced Analytics*, 607–613; <https://doi.org/10.1109/DSSAA49011.2020.00085>.
5. Culpepper, J.S., Faggioli, G., Ferro, N., and Kurland, O. Topic difficulty: Collection and query formulation effects. *ACM Trans. Inf. Syst.* 40, 1, Article 19 (Sept. 2021); <https://doi.org/10.1145/3470563>.
6. Demartini, G., Mizzaro, S., and Spina, D. Human-in-the-loop artificial intelligence for fighting online misinformation: Challenges and opportunities. *IEEE Data Eng. Bull.* 43, 3 (2020), 65–74; <http://sites.computer.org/debull/A20sept/p65.pdf>.
7. Pariser, E. *The Filter Bubble: What the Internet Is Hiding from You*. Penguin, U.K. (2011).
8. Saling, L.L. et al. No one is immune to misinformation: An investigation of misinformation sharing by subscribers to a fact-checking newsletter. *PLOS ONE* 16, 8 (August 2021), 1–13; <https://doi.org/10.1371/journal.pone.0255702>.
9. Thomson, T.J. et al. Visual mis/disinformation in journalism and public communications: Current verification practices, challenges, and future opportunities. *Journalism Practice* 16, 5 (2022), 938–962; <https://doi.org/10.1080/17512786.2020.1832139>.
10. Yom-Tov, E., Dumais, S. and Guo, Q. Promoting civil discourse through search engine diversity. *Social Science Computer Review* 32, 2 (2014), 145–154; <https://doi.org/10.1177/0894439313506838>

Damiano Spina is a senior lecturer of data science and DECRA Fellow at RMIT University, Melbourne, Australia.

Mark Sanderson is a professor of information retrieval and Dean of Research & Innovation at RMIT University, Melbourne, Australia.


Daniel Angus is a professor of digital communication at Queensland University of Technology, Brisbane, Australia.

Gianluca Demartini is an associate professor of data science at The University of Queensland, Brisbane, Australia.

Dana McKay is a senior lecturer of data science at RMIT University, Melbourne, Australia.

Lauren L. Saling is a senior lecturer of psychology at RMIT University, Melbourne, Australia.

Ryen W. White is a general manager and deputy lab director at Microsoft Research, Redmond, WA, USA.



Computing professionals must continue collaborating with other disciplines to make and integrate advances in critical areas, such as fairness, accountability, transparency, explainability, and the safety of human-AI cooperation.

