

Navigating Controversy as a Complex Search Task

Shiri Dori-Hacohen
Center for Intelligent
Information Retrieval
University of Massachusetts
Amherst
shiri@cs.umass.edu

Elad Yom-Tov
Microsoft Research
eladyt@microsoft.com

James Allan
Center for Intelligent
Information Retrieval
University of Massachusetts
Amherst
allan@cs.umass.edu

ABSTRACT

Seeking information on a controversial topic is often a complex task, for both the user and the search engine. There are multiple subtleties involved with information seeking on controversial topics. Here we discuss some of the challenges in addressing these complex tasks, describing the spectrum between cases where there is a clear “right” answer, through fact disputes and moral debates, and discuss cases where search queries have a measurable effect on the well-being of people. We briefly survey the current state of the art, and the many open questions remaining, including both technical challenges and the possible ethical implications for search engine algorithms.

1. INTRODUCTION

With the rise of personalization and the fear that it is creating a “Filter Bubble”, that is, exposure to a narrower range of viewpoints [31], navigating controversy is becoming an increasingly challenging task for search engine users and administrators alike. On one hand, by presenting answers to a user’s information need [9], search engines feed into confirmation bias and assist users - sometimes unawares - to remain in their own echo chambers. On the other hand, highlighting a controversy outright may have unintended consequences. The subtle differences between fact disputes and their interpretations, between scientific debates and moral stands, further exacerbate these challenges.

Information has a clear effect on the choices people make. The introduction of Fox News, a channel with clear political leanings, was associated with a shift of 3-8% in voting patterns in presidential elections from 1996 to 2000 towards the channel’s opinions [10]. In the health domain, queries about celebrities perceived as anorexic were shown to induce queries indicative of eating disorders [45].

Therefore, when a user’s information need pertains to a controversial topic, their search task becomes complex, as does the process of presenting the “correct” information. Since search engines match keywords to the retrieved docu-

ments, users are often left on their own to find the language used to describe different stances of an argument, in order to issue queries to retrieve information about them, and to classify the returned documents into these different views. Should search engines help users explicitly in this process? Should search engines make users aware of the different aspects of a topic or, alternatively, downweight some views (though this may arguably be viewed as censorship)? One way or another, helping the user navigate the controversial topic, along with its different opinions and stances, is a crucial part of the search engine’s role in the case of these complex search tasks, be it implicitly or explicitly.

Some might argue that the search engine’s role in the case of controversial topics ends at presenting the results in a simple keyword-based “list of ten links” on a Search Engine Results Page (SERP), and that the search engine has no place to take a moral stand. Even presenting the controversy and the various stances on it may not be a simple choice: if search engines provide comprehensive information on the different stances regarding a topic (e.g. presenting pro-anorexia opinions alongside anorexia treatments), this information may nudge people towards harmful behavior, either by exposing them to wrong or harmful information, or because users may stop perceiving search engines as honest brokers of information.

At the same time, simply providing every result available with no qualification can also be harmful, as disputed claims are allowed to proliferate without any warning to the unsuspecting user. For example, unproven, “quack” medical treatments often put users at risk by warning them not to heed their doctors [1, 4]. With unfounded claims widespread on the web, there are subtle ethical concerns with settling for a “buyer beware” (“caveat emptor”) approach. Caplan and Levin raise a similar concern regarding “caveat emptor” in the medical realm: “...researchers have an obligation to do more [in order] to enable patients to make informed choices” [8]. With concerns of life and death on the balance (e.g., in the case of medical controversies), we should not underestimate the impact of such choices on search engine users. Recent work assumes that trustworthiness should be preserved, for example in the case of knowledge extraction [11]. Some may go as far as arguing that, if technology allows for discernment of trustworthy vs. non-trustworthy sources, the search engine has an obligation to serve the trustworthy results to the users; others may say this is a slippery slope, and may in fact be viewed as censorship.

When discussing “navigating controversy” as a complex search task, there is an additional layer of complexity: be-

yond the complex task that the user herself is trying to complete, complexity also stems from the search engine’s design and algorithmic choices. It’s possible that amidst all the websites crawled by an engine, the correct response (if one even exists) is nowhere to be found, or is unfairly biased [42]. Should a search engine operator be concerned with civic or ethical implications of the search results it serves on controversial topics [20]? Should the user always be provided with what they want to see, even if it can be harmful to the user, or to society as a whole? Where should we draw the line between presenting trustworthy information from authoritative sources and discounting incorrect statements, versus presenting opinions on a moral debate?

These questions are open problems. Far from providing the community with a “correct” answer, we’d like to open the discussion on the case of navigating controversy as a complex search task. Here we highlight some of the issues that users may want to perform when searching for information on controversial topics, including seeking information on controversial topics; understanding different stances or opinions on such topics; and placing results within the context of the larger debate. Even the definition of controversy is an open question, which we will discuss as well.

2. SUPPORTING USERS WITH CONTROVERSIAL QUERIES

In order to account for users’ information needs on controversial queries and modulate the results in some way, there is first a technical challenge of recognizing that the query addresses a controversial topic, and determining what is controversial about it. Prior work has shown that it is possible to create classifiers for controversial Wikipedia pages [23, 35, 37] as well as events on Twitter [32]; recently, Dori-Hacohen and Allan demonstrated such a classifier to detect controversial web pages [12, 13].

Controversies can also be detected from a query perspective, if those are available, by finding queries that have semantically opposite meanings [19, 43]. Additionally, some advances have been made in recent years with regards to automatically detecting bias (cf. [33]). The goal of such detection could be to inform the user of the controversy by means of a browser extension or search engine warning [12]. A similar approach was demonstrated with regards to fact disputes [15], a specific type of controversy.

Assuming one has successfully discovered that a topic is controversial, another challenge is understanding what is controversial about it. In the political sphere, Awadallah et al. demonstrated automatic extraction of politician opinions [3]. Sentiment-based diversification of search on controversial topics has been proposed by Kacimi and Gamper [22], though several researchers have argued that controversy is distinct from sentiment analysis [3, 12, 27].

While frameworks for machine-readable argumentation and “The Argument Web” have been implemented [5], search engines cannot rely on widespread adoption of such tools. Recently, Borra et al. [6] demonstrated an algorithm that detects which topics are most contested within a given Wikipedia page; these and similar advances will be needed in order to present users with explicit stances on controversial topics.

3. THE PROBLEM OF DEFINING CONTROVERSY

How does one define controversy? While there is no one definition of the term controversy, we might use the following definition as an approximation: controversial topics are those that generate strong disagreement among large groups of people. Like the definition of relevance, it’s possible that controversy should be defined operationally: whatever people perceive as controversial, is controversial.

However, in line with others’ findings [24], our research so far shows that achieving inter-annotator agreement on the “controversy” label is very challenging. Additionally, while intuition and some researchers might suggest that the notion of sentiment should be relevant for controversy (e.g. [32, 39]), others have argued that sentiment is not the right metric by which to measure controversy [3, 12, 27]; opinions on movies and products may contain sentiment, yet lack controversy.

Likewise, we find some of the definitions of controversy used by others, or the datasets that those definitions lead them to use, to be very problematic (e.g. definitions that confound vandalism and controversy and therefore rate “podcast” as the most controversial topic in Wikipedia [40], or relying on the list of Lamest Edit Wars in Wikipedia as a controversy dataset [7]).

It may perhaps be helpful to break the definition of controversy into several interrelated definitions: For example, bias, disputes, truth value and polarity or intensity of emotion are potentially easier terms to define, but each of them only partially covers controversy. How does controversy relate to these constructs, and how would one proceed to discover the relationships between them?

Additionally, does the scope and context of the controversy matter? For example, do controversies regarding occurrences on American Idol (which may induce edit wars on Wikipedia) matter less than a controversy on the Israeli-Palestinian Conflict? One could argue that the latter is a much more controversial and influential topic; but for the user searching for “American Idol” or, for example, “Joanna Pacitti” (a controversial contestant on the show), perhaps the knowledge that this represents a controversial topic may be just as relevant – in the context of that search.

Though one may have an intuitive understanding of the term “controversy”, without a structured definition, our work (as well as others’) will not hold as much weight or predictive power.

3.1 Single truth or shades of gray

Information needs vary in the number of answers to them, both correct and incorrect. Some information needs have a single correct answer to them, while others may have several possible correct answers, requiring a moral judgment or entailing an opinion, e.g. political and religious questions. There are also questions for which there is a single scientifically correct answer, but for which non-scientific responses exist, even though they are factually incorrect. For example, some people claim that the Mumps-Measles-Rubella (MMR) vaccine causes autism; though studies have shown this claim to be incorrect, it is still believed by many people.

This variation in answers requires different treatment in each case. The simplest category is that where the information need has a single, correct, answer, which the search engine can provide. The second category is of questions

which have a technically correct response, but also an incorrect one which is prevalent on the web. Recent research by White and Hassan has demonstrated this phenomena in web search results, and specifically in health search [42].

The last category is of questions which have several possible correct answers, among which people may choose by making a moral judgment, for example, topics of abortion, same-sex marriage, and other highly charged issues; religious and political questions often fall under this umbrella. Selective exposure theory shows that people seek information which affirms their viewpoint and avoid information which challenges it [16]. Exposure to differing viewpoints has been shown to be socially advantageous in reducing the likelihood of adopting polarized views [36] and increasing tolerance for people with other opinions [17]. These advantages have led some to argue that technology could be used to expose people to a broader variety of perspectives, for example by modifying the display of information to nudge people to becoming “open-minded deliberators” [17].

This reasoning has led researchers to try and inform people of the differing views on the topics which they are reading. Providing people with feedback as to how much (on average) their reading was biased towards one or another political opinion, had only a small effect on nudging people to read more diverse opinions [28]. Kriplean et al. [26] developed a system for people to explicitly construct and share pro/con lists for a political election in Washington state, but found that opinions did not significantly change after using the system. In another experiment, Oh et al. [30] found that people preferred search results which were clearly delineated as to their leaning. Recently, Yom-Tov et al. [43] showed that people would read opposite opinions to theirs if their language model was appropriately selected. Such an intervention had long-lasting effects on reducing selective exposure. Thus, it is technically possible to provide people with diverse opinions where they have sought only one, but there still remains the question of whether this should be the role of a search engine.

An additional concern is whether claiming that certain facts are “true” or “false” holds any objective meaning. The scope of this paper does not allow a deep dive into the philosophical questions of objectivism vs. moral relativism, and the constructs of objectivity, subjectivity and intersubjectivity¹. Nonetheless, we can still delineate a few obvious concerns: the choice of which facts are in dispute, or which topics are controversial, can vary significantly with the cultural and social setting in which these questions are evaluated. For example, a user in Israel and a user in Iran may have very different opinions about what holds “true”, and either may be offended if the others’ worldview was presented as a “fact”; what is fact to one is either highly controversial or simply false to the other, and vice versa. As another example, the research by White and Hassan cited above [42] assumes that the Western world’s view of medicine is the only correct one, but users in China may beg to differ. Is a topic therefore only controversial if a user (or culture) believes it to be so? Who, then, can decide when a topic is controversial? How can the system know that a user believes a topic is controversial, and should the system then respond differently than when a user accepts it as “fact”?

¹Dubois [14] provides an insightful exploration of these concepts with regard to stance taking.

4. OPEN QUESTIONS

Several researchers have claimed that search engines have significant political power [21]. In his book *Republic.com 2.0*, legal scholar Cass Sunstein argues that a purely consumer-based approach to Internet search is a major risk for democracy [38]. One of deliberative democracy’s basic tenets, he argues, is the ability to have a shared set of experiences, and to be exposed to arguments you disagree with. Search engines and social media are increasingly responsible for “Filter Bubbles”, wherein click-feedback and personalization lead users to only see what they want, serving to further increase confirmation bias [31]. While this may seem to match individual users’ preference, the net effect on society is potentially detrimental. Being exposed only to like-minded people in so-called “echo chambers” serves to increase polarization and reduce diversity [34]².

Contrary to the common wisdom, some evidence exists that online personalization has not increased the filter bubble [18]. That said, research has shown that exposing users to opposing opinions increases their interest in seeking diverse opinions, and their interest in news in general [43]. There have been suggestions to diversify search results based on sentiment [22], though others argue that presenting the opposite opinion would only help in some cases [2, 29]. Prior bias of people changes the results of a search query, even without personalization. For example, the results for the query “what are the advantages of the MMR vaccine?” are completely different from the results served for the query “what are the dangers of the MMR vaccine?”. Moreover, the way people interpret the same information is dependent on their bias, for example in the case of gun control [25] or bias towards vaccines [44]. Thus, if a user seeks information on “how does MMR cause autism?”, should a search engine inform the user of the truth, or just satisfy their information need? One possible solution includes highlighting disputed claims [15] or explicitly presenting opposing viewpoints [41], but the problem remains that the user may not trust sources that don’t match their existing worldview.

Since search engines (as well as their social media counterparts) are increasingly the dominant medium for seeking information and news, the question then becomes: should search engines reflect what is on the internet and match content to users to maximize their preference, regardless of its truth value, or any concerns about diversity of opinion? Where do we draw the line between fact disputes and moral debates? Should the controversial nature of a topic depend on the social and cultural setting in which it is being evaluated? Should the search engines have a civic duty, and in that case, who decides what that duty is?

There are multiple technical challenges remaining in classifying controversial topics and extracting the opinions about them. However, even if these technical challenges of detecting controversy and stances were solved, there remains the question of if, when and how to present these to the user, based on their information need. As we discussed, there are ethical concerns with a search engine taking action, but also with inaction. It remains to be seen if users would be interested in hearing opposing opinions, or whether interventions would be useful; and finally, it is unclear whether it is within

²We note that, despite our own biases, the values of democracy and diversity of opinion are also culturally predicated, and not necessarily applicable to all search engine users.

the search engine's purview (or even its duty) to intervene, and if so, how.

Acknowledgments

This work was supported in part by the Center for Intelligent Information Retrieval and in part by NSF grant #IIS-1217281. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

We thank Gonen Dori-Hacohen, Myung-ha Jang, Shankar Kumar, Kishore Papineni and David Wemhoener for fruitful conversations. Special thanks to the anonymous reviewers for their insightful comments.

5. REFERENCES

- [1] American Cancer Society. Metabolic Therapy, Mar. 2012. Retrieved from <http://www.cancer.org/treatment/treatmentsandsideeffects/complementaryandalternativemedicine/dietandnutrition/metabolic-therapy>
- [2] J. An, D. Quercia, and J. Crowcroft. Why individuals seek diverse opinions (or why they don't). *Proceedings of the 5th Annual ACM Web Science Conference on - WebSci '13*, pages 15–18, 2013.
- [3] R. Awadallah, M. Ramanath, and G. Weikum. Harmony and Dissonance: Organizing the People's Voices on Political Controversies. *WSDM*, pages 523–532, Feb. 2012.
- [4] S. Barrett and V. Herbert. Twenty-Six Ways to Spot Quacks and Vitamin Pushers, 2014. Retrieved from <http://www.quackwatch.org/01QuackeryRelatedTopics/spotquack.html>
- [5] F. Bex, M. Snaith, J. Lawrence, and C. Reed. ArguBlogging: An application for the Argument Web. *Web Semantics: Science, Services and Agents on the World Wide Web*, 25:9–15, Mar. 2014.
- [6] E. Borra, A. Kaltenbrunner, M. Mauri, U. Amsterdam, E. Weltevrede, D. Laniado, R. Rogers, P. Ciuccarelli, and G. Magni. Societal Controversies in Wikipedia Articles. *Proceedings CHI 2015*, pages 3–6, 2015.
- [7] S. Bykau, F. Korn, D. Srivastava, and Y. Velegrakis. Fine-Grained Controversy Detection in Wikipedia, 2015.
- [8] A. Caplan and B. Levine. Hope, hype and help: Ethically assessing the growing market in stem cell therapies. *Current*, 10(5):33–34, 2010.
- [9] D. Carmel, E. Yom-Tov, A. Darlow, and D. Pelleg. What makes a query difficult? In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 390–397. ACM, 2006.
- [10] S. DellaVigna and E. Kaplan. The Fox News effect: Media bias and voting. *The Quarterly Journal of Economics*, 122(3):1187–1234, 2007.
- [11] X. L. Dong, E. Gabrilovich, K. Murphy, V. Dang, I. Watts, W. Horn, C. Lugaresi, S. Sun, and W. Zhang. Knowledge-Based Trust: Estimating the Trustworthiness of Web Sources. *Arxiv preprint*, (Section 3), 2015.
- [12] S. Dori-Hacohen and J. Allan. Detecting controversy on the web. In *Proceedings of the 22nd ACM international conference on Conference on Information & Knowledge Management, CIKM '13*, pages 1845–1848, New York, NY, USA, 2013. ACM.
- [13] S. Dori-Hacohen and J. Allan. Automated Controversy Detection on the Web. In *ECIR'15, To Appear. Preprint available at: http://maroo.cs.umass.edu/pub/web/getpdf.php?id=1173*, 2015.
- [14] J. W. Du Bois. The stance triangle. *Stancetaking in discourse: Subjectivity, evaluation, interaction*, pages 139–182, 2007.
- [15] R. Ennals, B. Trushkowsky, and J. M. Agosta. Highlighting disputed claims on the web. In *Proceedings of the 19th international conference on World wide web - WWW '10, WWW '10*, page 341, New York, New York, USA, 2010. ACM Press.
- [16] D. Frey. Recent research on selective exposure to information. *Advances in experimental social psychology*, 19:41–80, 1986.
- [17] R. K. Garrett and P. Resnick. Resisting political fragmentation on the Internet. *Daedalus*, 140(4):108–120, 2011.
- [18] M. Gentzkow and J. M. Shapiro. Ideological segregation online and offline. *Quarterly Journal of Economics*, 126:1799–1839, 2011.
- [19] K. Gyllstrom and M.-F. M. Moens. Clash of the typings: finding controversies and children's topics within queries. In *Proceedings of the 33rd European conference on Advances in information retrieval, ECIR'11*, pages 80–91, Berlin, Heidelberg, 2011. Springer.
- [20] L. M. Hinman. Esse est indicato in Google: Ethical and political issues in search engines. *International Review of Information Ethics*, 3(6):19–25, 2005.
- [21] L. D. Intronza and H. Nissenbaum. Shaping the Web: Why the Politics of Search Engines Matters. *The Information Society*, 16(3):169–185, 2000.
- [22] M. Kacimi and J. Gamper. MOUNA: Mining Opinions to Unveil Neglected Arguments. In *Proceedings of the 21st ACM international conference on Information and knowledge management - CIKM '12*, page 2722, New York, New York, USA, Oct. 2012. ACM Press.
- [23] A. Kittur, B. Suh, B. A. Pendleton, E. H. Chi, L. Angeles, and P. Alto. He Says, She Says: Conflict and Coordination in Wikipedia. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '07*, pages 453–462, New York, NY, USA, 2007. ACM Press.
- [24] M. Klenner, M. Amsler, and N. Hollenstein. Verb Polarity Frames: a New Resource and its Application in Target-specific Polarity Classification. In *Proceedings of the 12th edition of the KONVENS conference Vol. 1. - Hildesheim*. Universität Hildesheim, 2014.
- [25] D. Koutra, P. Bennett, and E. Horvitz. Events and Controversies: Influences of a Shocking News Event on Information Seeking. *TAIA workshop in SIGIR*, pages 0–3, 2014.
- [26] T. Kriplean, J. Morgan, D. Freelon, A. Borning, and L. Bennett. Supporting reflective public thought with ConsiderIt. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW)*,

- pages 265–274. ACM, 2012.
- [27] Y. Mejova, A. X. Zhang, N. Diakopoulos, and C. Castillo. Controversy and Sentiment in Online News. Sept. 2014.
- [28] S. A. Munson, S. Y. Lee, and P. Resnick. Encouraging Reading of Diverse Political Viewpoints with a Browser Widget. In *Proceedings of the International Conference on Weblogs and Social Media*, 2013.
- [29] S. A. Munson and P. Resnick. Presenting Diverse Political Opinions: How and How Much. In *Proc. CHI 2010*, CHI '10, pages 1457–1466, New York, NY, USA, 2010. ACM.
- [30] A. Oh, H. Lee, and Y. Kim. User evaluation of a system for classifying and displaying political viewpoints of weblogs. *Proc. ICWSM*, 2009.
- [31] E. Pariser. *The Filter Bubble: What the Internet is hiding from you*. Penguin Press HC, 2011.
- [32] A.-M. Popescu and M. Pennacchiotti. Detecting controversial events from twitter. In *Proceedings of the 19th ACM international conference on Information and knowledge management - CIKM '10*, pages 1873–1876, 2010.
- [33] M. Recasens, C. Danescu-Niculescu-Mizil, and D. Jurafsky. Linguistic Models for Analyzing and Detecting Biased Language. *Proceedings of the 51st Annual Meeting on Association for Computational Linguistics*, 2013.
- [34] D. Schkade, C. R. Sunstein, and R. Hastie. What happened on deliberation day? *California Law Review*, 95(298):915–940, 2007.
- [35] H. Sepehri Rad and D. Barbosa. Identifying Controversial Articles in {Wikipedia}: A Comparative Study. In *Proceedings of 8th conference on WikiSym*, WikiSym '12. ACM, 2012.
- [36] A. L. Stinchcombe. Going to Extremes: How Like Minds Unite and Divide. *Contemporary Sociology: A Journal of Reviews*, 39(2):205–206, 2010.
- [37] R. R. Sumi, T. Yasseri, A. Rung, A. Kornai, and J. Kertész. Edit wars in Wikipedia. *Privacy, Security, Risk and Trust (PASSAT), 2011 IEEE Third International Conference on Social Computing (SocialCom)*, pages 724–727, 2011.
- [38] C. R. Sunstein. *Republic.com 2.0*. Princeton University Press, 2009.
- [39] M. Tsytsarau, T. Palpanas, and K. Denecke. Scalable detection of sentiment-based contradictions. *DiversiWeb 2011*, 2011.
- [40] B.-q. Vuong, E.-p. Lim, A. Sun, M.-T. Le, H. W. Lauw, and K. Chang. On ranking controversies in Wikipedia: models and evaluation. In *Proceedings of the international conference on Web search and web data mining*, WSDM '08, pages 171–182, New York, NY, USA, 2008. ACM.
- [41] V. G. V. Vydiswaran, C. Zhai, D. Roth, and P. Pirolli. BiasTrust: Teaching Biased Users About Controversial Topics. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management*, CIKM '12, pages 1905–1909, New York, NY, USA, 2012. ACM.
- [42] R. W. White and A. Hassan. Content bias in online health search. *ACM Transactions on the Web (TWEB)*, 8(4):25, 2014.
- [43] E. Yom-Tov, S. T. Dumais, and Q. Guo. Promoting civil discourse through search engine diversity. *Social Science Computer Review*, 2013.
- [44] E. Yom-Tov, L. Fernandez-Luque, and L. Luque. Information is in the eye of the beholder: Seeking information on the {MMR} vaccine through an Internet search engine. In *Proceedings of the American Medical Informatics Association*, 2014.
- [45] E. Yom-Tov and D. m. boyd. On the link between media coverage of anorexia and pro-anorexic practices on the web. *International Journal of Eating Disorders*, 47(2):196–202, 2014.