BiasTrust: Teaching Biased Users About Controversial Topics

V.G.Vinod Vydiswaran† ChengXiang Zhai† Dan Roth† Peter Pirolli‡
†University of Illinois, Urbana, IL
‡Palo Alto Research Center, Palo Alto, CA
†{vgvinodv,czhai,danr}@illinois.edu, ‡Peter.Pirolli@parc.com

ABSTRACT

Deciding whether a claim is true or false often requires understanding the evidence supporting and contradicting the claim. However, when learning about a controversial claim, human biases and viewpoints may affect which evidence documents are considered "trustworthy" or credible. It is important to overcome this bias and know both viewpoints to get a balanced perspective. In this paper, we study various factors that affect learning about the truthfulness of controversial claims. We designed a user study to understand the impact of these factors. Specifically, we studied the impact of presenting evidence with contrasting viewpoints and source expertise rating on how users accessed the evidence documents. This would help us optimize how to teach users about controversial topics in the most effective way, and to design better claim verification systems. We find that users do not seek contrasting viewpoints by themselves, but explicitly presenting contrasting evidence helps them get a well-rounded understanding of the topic. Furthermore, explicit knowledge of the source credibility and the context not only affects what users read, but also how credible they perceive the document to be.

Categories and Subject Descriptors

H.1.2 [**Information Systems**]: Models and Principles—*Human information processing, Human factors*

Keywords

Claim verification, Information credibility, User study

1. INTRODUCTION

The Web today is a hodgepodge of well-curated, edited content and freelance, unmoderated content. Typically, well-structured, formatted, and edited content is considered more trustworthy and credible, while information that appears in forums and message boards is considered not as trustworthy. But history is replete with many instances where credible sources have made significant errors in stating facts or helped spread rumors; possibly due to their own biases. This indicates a growing challenge for users to gather an unbiased opinion about topics of interest over the Internet and also provides a strong motivation to study it.

Consider a scenario where an Internet surfer wants to learn about a recent controversy. To give an example, consider that Alice wants to know if chocolate and flavored milk provided in schools is a healthy food choice for her kids. Depending on the keywords she chooses to search about the topic, she might see news articles about a recent ban on chocolate milk in certain schools, or learn about health benefits of milk in growing kids. She might find results from news media organizations reporting on the ban, or activist groups actively encouraging drinking milk, or even concerned parents posting questions and responding to them via community-driven question answering services.

Is Alice equally likely to read these articles, and read them in the order presented? The answers appear to be in the negative, but we wanted to understand what factors impact these decisions. Specifically, we wanted to study the factors that influenced (i) the documents users read, (ii) the extent of learning, and (iii) the perceived credibility of a source. We believe this is one of the first works that studies these aspects, especially when learning about controversial topics.

We designed and conducted a user study to understand how to present evidence documents to enable users to learn about the controversy in an unbiased fashion and help them overcome any prior bias they might have about the topic. First, we try to understand the user's position and beliefs about the controversy. Then, we study if the display of contrasting viewpoints to evidence documents and source expertise ratings influence user behavior. We find that users tend not to seek contrasting viewpoints, at least in a limited-time learning scenario, but that explicitly presenting contrasting evidence helps them get a balanced understanding of the topic. Furthermore, we observe that explicit knowledge of the source credibility or expertise, and the context in which the evidence was provided, not only affects what users read, but also how credible they perceive the documents to be. These insights help us optimize the presentation of credible evidence documents to teach biased users about controversial topics in the most effective way.

In this short paper, we explain the design of the user study and the interface variants in detail. However, due to space constraints, we restrict our analysis to factors that help users get a balanced understanding of the topic. Additional findings are presented in Vydiswaran et al. [9].

CIKM'12, October 29–November 2, 2012, Maui, HI, USA.

2. RELATED WORK

Understanding what documents people read is related to research in many fields. Psychologists have studied the phenomenon of confirmation bias [5, 1], which states that people tend to favor information that confirms their beliefs – not only in deciding what to read, but also how they interpret what is read. Similarly, researchers in political science [7] found that people process information in biased fashion. They uncritically accept supporting arguments and argue against those that are contrary to their beliefs.

Researchers have looked at aggregating information from multiple sources to answer specific questions. Wu and Marian [10] studied how to collect information from multiple sources over the Web, specifically to find answers using corroborative evidence from multiple sources. Gallard et al. [2] looked at collecting information from contrasting viewpoints. The work presented in this paper looks at the next step of how to present these documents with contrasting viewpoints to users, while providing information about credibility of the sources, to enable users to learn about the topic efficiently.

Researchers have also looked at factors that influence what information users access and how they process them. This includes work on building tools to increase transparency and credibility of Wikipedia articles [4, 6] to help users decide if the information is credible or not. Ugander et al. [8] studied how to convince people based on the influence of one's social network on their actions. Extending it to the information domain, this research suggests that users should be exposed to multiple viewpoints. Pariser [3] investigated the notion of *filter bubble*, where search engines personalize web search results to show users information similar to what they have already seen. This not only encourages confirmation bias, but also excludes contradictory viewpoints. Our study tries to understand how to overcome these shortcomings by presenting contrasting, yet credible evidence to users. We believe our work is the first to study how source expertise rating, evidence context, and contrasting viewpoints help users learn about controversial topics.

3. BIASTRUST: DESIGNING THE STUDY

Understanding which claims to believe and why to believe those claims is important to make an informed decision. This is basically a learning task, where an inquisitive user tries to learn as much as possible about the claim and assimilate all evidence in support of or against the claim. It is an interesting challenge for a retrieval system to not only retrieve documents relevant to the claim, but also present it succinctly to help users understand that information quickly. However, as we pointed out in Sec. 2, previous research by psychologists and others have shown that users tend to access information that supports their own viewpoints. So, it is important for an automated claim verification system to model such human biases and present trustworthy evidence to overcome this bias, where possible.

We designed a system that would retrieve relevant, trustworthy documents about a topic and provide the user with an overall, unbiased perspective. We conducted a user study to learn how to optimize the display of credible information to users. The BiasTrust study investigates the factors that may influence humans to decide what to read and what to treat as credible. We focused on three major factors, viz. (i) the ability to access credible, yet contrasting viewpoints about a claim to help gauge the trustworthiness of the claim; (ii) the knowledge of source expertise and credibility, and how that affects the decision on what evidence is read; and (iii) the order in which the documents are presented, and if that affects the overall understanding of the topic. Other factors, such as summarizing documents and assigning an overall truth value to a claim may also be relevant in designing a claim verification system. However, we chose not to study these factors, but instead provide access to the relevant evidence to allow users verify claims by themselves. Further, the study focused on controversial claims instead of factual claims. Controversial claims have many genuine evidence documents supporting and opposing the claims, so by focusing on such claims, we could study how preference-based factors affect the learning process.

The user study was designed as a learning task, where users were asked to learn as much as possible about a topic within a stipulated time. The study was conducted in three online stages, viz. (i) Pre-study survey questionnaire, (ii) the Study phase, and (iii) Post-study questionnaire.

Stage 1: Pre-study questionnaire: In the pre-study survey, subjects were asked questions that helped us gauge their (lack of) knowledge and bias towards/against issues relevant to the topic being studied. Subjects answered the questions on a four-point Likert scale. For the knowledge questions, the response scale ranged from (i) 'insignificant' to (iv) 'very significant', while for bias questions, it ranged from (i) 'strongly against' to (iv) 'strongly in favor of' the issue. Subjects could also respond to a question with an 'I don't know' answer. This design of using knowledge and bias related questions and limiting the nature of allowed responses was to encourage subjects to think about their position on many sub-topics related to the overall issue.

Stage 2: Study phase: Once the responses to the pre-survey questionnaire were recorded, subjects were directed to one of the interface variants (described in detail in Sec. 4). In each interface variant, subjects had some contextual information about the passages. For instance, subjects were shown the source of the passage, its sub-topic, and whether the passage was in favor or against the sub-topic. For each passage subjects read, they were asked to answer two questions about the passage, viz. whether they agreed with what was being said in the passage and if they believed the information was biased with respect to the topic. Subjects were asked to continue reading until they believed they had read enough about the topic. Once they decided to quit the study phase, they were taken to the third and final stage of the online study.

Stage 3: Post-study questionnaire: After subjects spent time learning about the topic in the study phase, they were asked to answer the same set of topic-specific questions that were posed during the pre-study questionnaire. Subjects responded based on the topics they read about and how important and relevant they felt the sub-topics were, after the study. They were also asked to provide feedback on what interface features helped them in their task, and suggest additional features that might help them.

4. USER INTERFACE VARIANTS

We designed interfaces to study the following factors: (i) explicit display of contrasting evidence, (ii) single document per page vs. multiple documents, (iii) source expertise rating, and (iv) presentation order of documents.

Figure 1: Single document view with option to look at contrasting document (UI# {$1a$, $1b$})

Figure 2: Single document view that shows contrasting document by default (UI# {$2a$, $2b$})

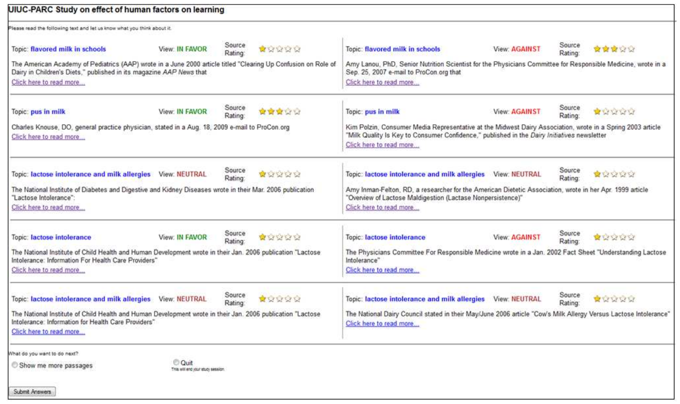

Figure 3: Multi-document view that shows five primary documents, along with the corresponding contrasting documents (UI# {3, $4a$, $4b$, 5})

Variant identifier	# passages	contrast view	topics sorted	source rating	
				show	scheme
UI# $1a$	1	No	No	Yes	Bimodal
UI# $1b$	1	No	No	Yes	Uniform
UI# $2a$	2	Yes	No	Yes	Bimodal
UI# $2b$	2	Yes	No	Yes	Uniform
UI# 3	10	Yes	No	Yes	Bimodal
UI# $4a$	10	Yes	Yes	Yes	Bimodal
UI# $4b$	10	Yes	Yes	Yes	Uniform
UI# 5	10	Yes	Yes	No	—

Table 1: Parameter settings for interface variants

1. Explicit display of contrasting evidence: We believe that one of the major factors in learning about a controversial topic in an unbiased fashion is the exposure to alternate viewpoints. To verify if this conjecture is true, we designed two variants of the system. In the first variant, users were exposed to just one document from a single viewpoint at a time, in a fixed order. Users could, however, explicitly ask for the next document to be one of opposing viewpoint. If not, the next relevant result was shown by default. UI variants $1a$ and $1b$ (cf. Table 1) follow this setting, and Fig. 1 shows an example interface.

In the second variant, users were exposed to the contrasting evidence right at the start. The primary document and a document of contrasting viewpoint were shown side-by-side. Users could still pick which document to read first, and may choose to ignore the contrasting viewpoint if they wish to. UI variants $2a$ and $2b$ (cf. Table 1) follow this setting, and Fig. 2 shows an example.

2. Single document per page vs. multiple documents: The next factor we investigated was how the amount of information on a page affected the reading pattern and choice of documents. When fewer documents are shown, humans tend to spend more time reading documents that are shown, before moving on to the next page. We wanted to investigate if this hypothesis was correct. We designed a third variant where instead of one document, five documents and their corresponding counter-arguments were shown to the user on a single page. UI variants {3, $4a$, $4b$, 5} show multiple documents per page, and Fig. 3 shows an example.

3. Source expertise rating: Another factor in deciding what to read is whether users believe the document comes from a credible source. To test this hypothesis, we decided to control the expertise ratings in two ways. In one UI variant (UI# 5), the expertise ratings were hidden, and subjects did not know if the source was credible or not.

The second way we controlled the expertise rating was to show two different rating schemes. In one scheme, all sources got a rating of either 1-star or 3-star. We call this the *bimodal rating scheme*, since the ratings appear to be drawn from a bimodal distribution. In the second scheme, sources were assigned a random rating drawn from a uniform distribution while ensuring that the rating was not the same as the one under the bimodal scheme. We call this the *uniform rating scheme*. Sources were assigned this rating in a static fashion; that is all participants that were shown one rating scheme, consistently saw the same expertise rating for the same source. The UI variants {$1a$, $2a$, 3, $4a$} followed the bimodal rating scheme, while the three remaining variants {$1b$, $2b$, $4b$} followed the uniform rating scheme.

4. Document presentation order: The final hypothesis we wanted to test was whether grouping documents together by sub-topic significantly helped the learning task and affected the selection of documents read. We focused on testing this hypothesis only in the case when multiple documents are shown. We had two configuration settings: (i) passages are shown in relevance order, and hence appear to come in a random sequence of sub-topics (in UI variants {$1a$, $1b$, $2a$, $2b$, 3}); and (ii) passages are shown in topic-sorted order, in which topics and passages from a topic are both ordered by relevance (in UI variants {$4a$, $4b$, 5}).

5. SETTING UP THE USER STUDY

Data and Study topics: We enabled the study for two controversial issues, one from the health domain and the other from politics. The primary claims ("*issues at hand*") included in this study were: (i) `Milk`: Drinking milk is a healthy choice for humans; and (ii) `Energy`: Alternate sources of energy are viable alternatives to fossil fuels. We will refer to the corresponding study sessions as the `Milk` and `Energy` tasks, respectively.

For both issues, we collected over 350 snippets of text from ProCon.org[1], a non-partisan, non-profit public charity website. The website offers quotes relevant to sub-topics within the issue, categorized as `pro` (in favor of the question being asked), `con` (against the question being asked), or neither `pro` nor `con`. The website also gives an expertise rating on a five-point scale to each source. Our analysis showed that for both issues, almost all sources were assigned either a 3-star or a 1-star rating. We used these manually assigned ratings as our bimodal rating scheme.

We indexed the passages using the Lemur toolkit[2]. We fixed the retrieval query as a weighted combination of key words relevant to the topics, and retrieved top 50 passages. This helped us control the exact set and sequence of passages that subjects would see in the user study.

Inviting subjects: Volunteers were invited to participate in the study by announcing the study on mailing lists in many departments within a large, diverse, public university. Invitation emails were sent out primarily to graduate students, staff members, and to participants of a nation-wide multi-center research meeting. The announcement invited volunteers to participate in a learning task, where they were expected to learn about a topic and answer questions. Subjects were not informed about the exact nature of the study or what factors were being measured. They were, however, informed that they can participate in the learning tasks related to two topics, that each task would take about 45 minutes to complete, and that they will be rewarded a fixed amount for each task they successfully complete.

Volunteers who agreed to participate in the study were issued a unique identifier (pseudonym) that they would use to access the study. The pseudonyms were statically mapped to two interface variants from the ones listed in Table 1. One task was assigned a variant from the first four that showed one or two documents per page, and the other task was assigned a variant from the latter four that showed ten documents per page. Subjects were free to choose any of the two tasks first, so the researchers did not have control over which interface variant subjects saw, and for subjects that took part in both topics, the order in which they saw those variants. All responses and interactions were recorded using the pseudonym and no identifiable information was requested or recorded during the study.

6. ANALYSIS OF STUDY RESULTS

We invited 24 volunteers to participate in the user study, and the average age of participants was 28.6 ± 4.9 years. Each participant could take part in at most two study tasks. In all, we collected information from 40 study tasks, with most participants choosing to take part in both tasks. The profile of how subjects interacted with the system was similar for both tasks. Typically, subjects took 7–10 minutes to complete each of the pre-study and post-study questionnaires, and spent about 26.5 minutes in the study phase.

6.1 Which documents do subjects read/skip?

In this paper, we focus on the factors that help users get a balanced understanding of the topic. Specifically, we want to understand how contrasting viewpoints and expertise ratings help users decide what documents to read.

[1] `http://www.procon.org/`

[2] `http://www.lemurproject.org`

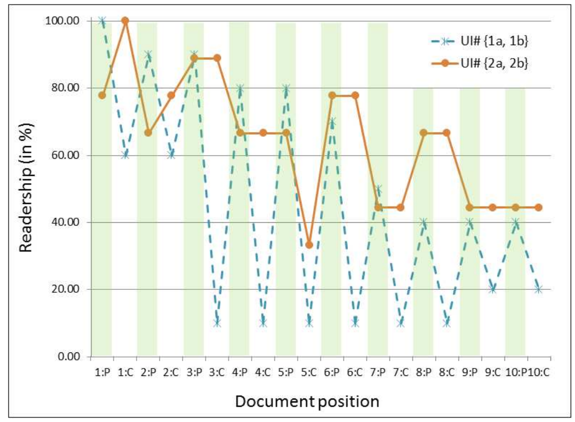

Figure 4: Variation of number of documents read for each document position. Documents in the primary set are in odd positions (shaded green) while those in the contrast set are in even positions (shaded white).

(i) Explicitly showing contrast documents helps

Subjects read an average of 18.6 documents, but considered as many as 31 documents, including ones they choose to skip. To understand which documents subjects read and which they skip, we compared how many times subjects read a document per result position. In UI# $1a$ and $1b$, subjects are shown only one document by default, from the primary document set; and they can choose to see the corresponding contrast document next. For all other UI variants, the contrast documents are shown along side the primary documents, with documents in favor of the sub-topic on the left column, and documents against the sub-topic on the right.

Fig. 4 shows how the readership changes with document position for the top 10 results. We compare two scenarios – one in which only one primary document is shown by default (UI# $\{1a, 1b\}$) and the other where one primary and one contrast document are shown side-by-side (UI# $\{2a, 2b\}$). As we see, the readership for contrast documents is significantly lower when it is not shown by default, while the readership of the primary document does not change as much. This shows that users do not tend to pro-actively ask for the contrasting viewpoint. But when shown side-by-side, they tend to read both viewpoints.

(ii) Showing multiple documents per page helps

When multiple documents are shown per page, users tend to be more selective in what they read. Table 2 summarizes the variation in reading pattern as number of documents are increased. Showing one document per page not only significantly reduce the total number of documents read, but subjects also spend more time reading those documents on-an-average. By showing the contrast document along side, subjects tend to spend more time overall to read more documents. When multiple documents are shown on a page, subjects are able to consider and skip many more documents in the stipulated time. When we compare single document and multi-document views, we see that although subjects read the same number of documents overall, they scanned about 15 more passages in the multiple document view.

UI# #passages	{1a, 1b} 1	{2a, 2b} 2	{3, 4a, 4b, 5} 10
Number of docs read	13.5	20.3	20.2
Number of docs skipped	4.3	5.1	19.7
Study time (in min)	26.3	29.4	25.4
Time/doc read (in min)	2.38	1.77	1.43

Table 2: Interface type influences reading pattern.

Scheme	**1★**	**2★**	**3★**	**4★**	**5★**
Single document view (UI# {1a, 1b, 2a, 2b})					
Bimodal	76.3%		85.4%		
Uniform	56.3%	77.3%	85.3%	83.2%	84.3%
Multiple document view (UI# {3, 4a, 4b})					
Bimodal	39.7%		61.3%		
Uniform	38.2%	52.8%	60.7%	66.0%	66.9%

Table 3: Readership increases with expertise rating.

(iii) Higher expertise rating gets higher readership

Next, we looked at the impact of expertise rating on which documents are read and which are skipped. For this analysis, we distinguish the UI variants based on whether the rating scheme was bimodal or random. Table 3 shows the variation in documents read for these two rating schemes.

We focus on two classes of interfaces – one in which a single primary document (and the corresponding contrast document) was shown, and the other where five primary documents were shown. In both cases, we find that the documents that are rated high are read more often than those with poor expertise ratings. This is especially true when multiple documents are shown and subjects have more choice in what to read. As we see in Table 3 (last row), only 38% of documents with a 1-star rating were read, whereas about 67% of documents with a 5-star rating were read by the subjects. When the documents are rated under the bimodal scheme, the documents rated high are 22% more likely to be read than the documents rated low. These trends also follow when only one document (with or without the contrast document) is shown to the subjects. We also see that, in this scheme, relatively fewer documents are skipped. This is because, in the single document view, there is typically just one option – i.e. to either read or skip the document, without the knowledge of what the next document would be. So, a relatively larger number of documents are read, even if they have low expertise ratings.

(iv) Absence of rating helps low quality documents

In one of the UI variants (UI# 5), all expertise ratings were hidden from the subjects. We find that, under this setting, 49.8% of the shown documents were read by the subjects. Comparing to Table 3, we find that unrated documents are more likely to be read than those with 1-star rating.

7. CONCLUSION AND FUTURE WORK

Providing access to unbiased and credible information is critical to satisfy information need in many domains. Moreover, when verifying controversial claims, it is also important to understand and help users overcome the tendency to stick to one's own viewpoints. We conducted a user study called BiasTrust to understand what factors significantly help users learn about a controversial topic. We find that explicitly showing contrasting viewpoints is significantly better than merely giving an option to users to look at documents from alternate viewpoints. Moreover, showing expertise rating helps users pick which documents to read and which to omit. We also observe that these factors help users reduce strong biases and learn new topics in an unbiased fashion [9].

Although this is an initial study on how to present controversial topics to potentially biased users, the findings are already interesting. Subjects spent over an hour for each task and gave valuable feedback during the course of their participation. The insights gained by this study would help us and other researchers optimize the design of an automated claim verification system, that not only learns which evidence documents are most relevant to show to the user, but also what additional information needs to be provided to help users assimilate the information faster.

This study can be further extended to understand how human biases affect the perceived quality of the documents. Further, based on users' prior knowledge, a different set of results may be presented. A larger scale study is planned to look into these aspects. Some subjects suggested that an overall summary of relevant sub-topics would help them understand the issues better. This is indeed an interesting need, but beyond the scope of the current study.

8. ACKNOWLEDGMENTS

This research is partially funded by grants from the MIAS Center of Excellence of the Department of Homeland Security, the Information Trust Institute, the Army Research Laboratory, and the Office of Naval Research.

9. REFERENCES

[1] J. Baron. *Thinking and Deciding*. Cambridge Press, 2000.

[2] A. Galland, S. Abiteboul, A. Marian, and P. Senellart. Corroborating Information from Disagreeing Views. In *Proc. of WSDM*, pages 131–140, 2010.

[3] E. Pariser. *The Filter Bubble: What the Internet is Hiding from You*. Penguin Press, 2011.

[4] P. Pirolli, E. Wollny, and B. Suh. So You Know You're Getting the Best Possible Information: A Tool that Increases Wikipedia Credibility. In *Proc. of CHI*, pages 1505–1508, 2009.

[5] S. Plous. *The Psychology of Judgment and Decision Making*. McGraw-Hill, 1993.

[6] B. Suh, E. H. Chi, A. Kittur, and B. A. Pendleton. Lifting the Veil: Improving Accountability and Social Transparency in Wikipedia with WikiDashboard. In *Proc. of CHI*, pages 1037–1040, 2008.

[7] C. S. Taber and M. Lodge. Motivated Skepticism in the Evaluation of Political Beliefs. *American Journal of Political Science*, 50(3):755–769, July 2006.

[8] J. Ugander, L. Backstrom, C. Marlow, and J. Kleinberg. Structural Diversity in Social Contagion. In *Proceedings of the National Academy of Sciences of the United States of America*, pages 1–5, 2012.

[9] V. Vydiswaran, C. Zhai, D. Roth, and P. Pirolli. Unbiased Learning of Controversial Topics. In *Proc. of ASIST*, 2012.

[10] M. Wu and A. Marian. Corroborating Answers from Multiple Web Sources. In *WebDB*, pages 1–6, 2007.