Probabilistic Approaches to Controversy Detection

Myungha Jang, John Foley, Shiri Dori-Hacohen, and James Allan

Center for Intelligent Information Retrieval
College of Information and Computer Sciences
University of Massachusetts

{mhjang, jfoley, shiri, allan}@cs.umass.edu

ABSTRACT

Recently, the problem of automated controversy detection has attracted a lot of interest in the information retrieval community. Existing approaches to this problem have set forth a number of detection algorithms, but there has been little effort to model the probability of controversy in a document directly.

In this paper, we propose a probabilistic framework to detect controversy on the web, and investigate two models. We first recast a state-of-the-art controversy detection algorithm into a model in our framework. Based on insights from social science research, we also introduce a language modeling approach to this problem. We evaluate different methods of creating *controversy language models* based on a diverse set of public datasets including Wikipedia, Web and News corpora.

Our automatically derived language models show a significant relative improvement of 18% in AUC over prior work, and 23% over two manually curated lexicons.

1. INTRODUCTION

The power of search is inherently coupled to the trustworthiness and reliability of the dataset the user is searching. Web search can be both helpful and troubling in this regard. In medicine, one may find misinformation that professionals consider incorrect, whether it is fraudulent treatments or side-effects that simply do not exist, such as the link between vaccines and autism. In the political domain, increasing polarization of sources and viewpoints leads to “Filter Bubbles” [15], where users are only able to find results that agree with them or the way they formulated their query. In these domains and more, the ability to reliably detect the presence of controversy in search results could aid critical literacy and the ability to make real-life decisions. Our task is to determine if a given Web document discusses controversial topics.

However, existing search engines are unlikely to reveal controversial topics to users unless they already know about

the controversies [10]. There is an increasing call for search engines to detect these queries and address them appropriately [9, 12]. Previous work [8, 11] presented an algorithm for classifying controversy in Web documents using Wikipedia; this reliance introduces efficiency concerns, and limits the classifier’s potential scope to topics covered in Wikipedia. Additionally, to date, little effort has been made to formulate a probabilistic model to detect controversy.

In this work, we investigate probabilistic models for automated controversy detection and thus provide theoretic modeling of the problem. First, we develop a probabilistic framework for the controversy detection problem and re-cast the state-of-the-art algorithm in that probabilistic perspective. We use this new perspective to extend prior work from a binary classification model to a probabilistic model that can be used for ranking. We then introduce the *controversy language model*, a new approach to address the controversy detection problem. We develop this language-modeling approach based on a recent qualitative analysis of controversy [5], which is more efficient and easier to compute than the more expensive Wikipedia-based controversy detection algorithm. Finally, we empirically validate our controversy language models and find that our approach significantly outperforms the state-of-the-art algorithm for controversy detection [8].

2. RELATED WORK AND BACKGROUND

Controversy detection studies often focus on certain genres or data sources, such as news [3, 14], Twitter [16], and Wikipedia [13, 17, 21]. Since these studies frequently use data-source-specific features such as Wikipedia’s edit history features or Twitter’s social graph information, it seems unlikely that existing work can be easily generalized to controversy detection on arbitrary webpages.

While some past work uses sentiment as a signal when researching controversy [2, 3], others have argued that opinion and controversy are distinct and non-overlapping concepts [1]. Researchers have shown that using sentiment for controversy detection performs poorly on webpages [8] and that controversy and sentiment are not directly related [14].

Choi et al. attempt to identify controversy and controversial subtopics using various features, particularly a mixture model of topic and sentiment [3]. We depart from their work by directly modeling the probability of controversy.

There also have been a few attempts to detect controversial content with lexicons. Roitman et al. retrieve Wikipedia articles with claims about controversial query topics [18], and Mejova et al. use crowdsourcing to label controversial words [14].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM’16, October 24–28, 2016, Indianapolis, IN, USA

© 2016 ACM. ISBN 978-1-4503-4073-1/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2983323.2983911>

In the social sciences, Cramer explains that the term “controversy” tends to be metadiscursive. Writers label something as “the controversy” by using their writing to categorize others’ writings and speakings about a topic. Hence, the phrase “the controversy” is indexical of controversial events because it is a linguistic way of pointing at a particular event. Cramer’s work suggests that language could be a key feature in identifying controversy.

2.1 Wikipedia Controversy Features

Previous work studied algorithms that generate scores that signal controversy in Wikipedia [13, 7, 21]. We refer to these scores as Wikipedia Controversy Features (WCF). The scores use various information extracted from Wikipedia pages, meta-data, talk pages, and the page’s edit history.

We use the following features to reimplement the k NN-WC algorithm (see §2.2) which relies on them, as well as to collect highly controversial documents for our new language modeling approach:

D Score The presence of the `{dispute}` tag placed on an article by editors. Only 0.03% of articles contain this tag.

C Score is generated by a regression-based method [7, 13] that trains on revisions that are labeled with `{controversial}` tags.

M Score is generated by Yasseri et al. based on features of edits to predict the ferocity of an “edit war” [21].

2.2 k NN-WC Algorithm

Dori-Hacohen and Allan [8] presented an algorithm for identifying whether a given Web document discusses controversial topics. To the best of our knowledge, this work is the first and only attempt to extend the controversy detection problem to general web pages in open domain [8], and build upon past work that investigated controversy detection on Wikipedia. We will refer to their work as k Nearest Neighbors of Wikipedia Controversy (k NN-WC).

The k NN-WC algorithm assumed that the controversy in a web document can be detected from the levels of controversy of related topics. The algorithm modeled related topics as related Wikipedia articles, and existing controversy labels on the Wikipedia neighbors [13, 21] were used to derive a final judgment on the original document (see §2.1).

Overview of k NN-WC Algorithm: k NN-WC [8] consists of four steps: (1) Find k Wikipedia neighbors given a webpage, (2) compute the three WCF scores (D, C, and M) for each neighbor, (3) aggregate the k values of each WCF score and turn them into three binary labels using thresholds, and (4) vote the labels and make a final decision.

Limitations: k NN-WC is constrained by its dependency on Wikipedia controversy indicators, which are sourced solely from Wikipedia-specific features such as mutual reverts of editors. Search of k -NN for each document is a non-trivial operation, which raises practical efficiency issues. In addition, topical coverage in Wikipedia is necessitated.

In our work, we first contribute a generalization of this algorithm to a probabilistic framework, grounding it in a theoretical conception. We then depart from it by presenting a new language model for controversy.

3. PROBABILISTIC MODELING OF THE k NN-WC ALGORITHM

In this section, we formulate the controversy detection problem with a probabilistic perspective and recast the k NN-

WC algorithm as obtaining a probability of controversy for binary classification. We then extend the k NN-WC *algorithm* for binary classification and derive the k NN-WC *model*, a probabilistic model that can be used for ranking.

We formulate the controversy detection problem as obtaining the following probabilities: Let D be a document. We define $P(C|D)$ as the probability that D is controversial, and $P(NC|D)$ as the opposite (i.e., non-controversial), such that the two probabilities sum to 1.

$P(C|D)$ can be represented by two components: $P(C, D)$, a joint probability of the presence of controversy and D ; and $P(D)$, a prior probability of D . For binary classification, we are interested in whether $P(C|D) > P(NC|D)$. However, $P(D)$ is the same for $P(C|D)$ and $P(NC|D)$, and we can thus ignore it in comparison. We model the approach of the k NN-WC algorithm by assuming that $P(C|D)$ can be estimated from the joint probability $P(C, D)$ by the following equation. We then assume that a document and its topics, denoted as T_D are interchangeable.

$$P(C|D) = \frac{P(C, D)}{P(D)} \approx \frac{P(C, T_D)}{P(D)} \quad (1)$$

In the context of k NN-WC, we interpret Wikipedia neighbors as a set of “latent topics” in the document. We denote a set of similar Wikipedia articles to the document (i.e., neighbors) as W_D . W_D was obtained by matching the query generated from the D to the Wikipedia index.

$$P(C, T_D) \approx P(C, W_D) \quad (2)$$

By treating the Wikipedia neighbors as topics, we are able to model $P(C, W_D)$ via WCF (see §2.1). By construction, a WCF reflects some estimate of the co-occurrence of controversy and the topic.

Likewise, the k NN-WC algorithm used a binary indicator function $P(C, W_D) \approx c(W_D)$ that outputs 1 (controversial) or 0 (non-controversial). One of the best performing settings of this algorithm [8]:

$$c(W_D) = \mathbb{1}[\max(WCF_M(W_D)) > \theta_M] \quad (3)$$

where $WCF_M(W_D)$ is a set of M scores for W_D , θ_M is a threshold, and $\mathbb{1}$ converts true to 1 and false to 0.

To support further evaluation of the k NN-WC model, we extend the indicator function to a numeric scoring function by removing the threshold so that the model can produce a ranking instead of binary classes. Let $f(W_D)$ be a function that returns an aggregated value of the given neighbors’ WCF scores.

$$f(W_D) = \text{agg}(WCF_{\{M, C, D\}}(W_D)) \quad (4)$$

We then convert this aggregation function to a probability by normalizing over all possible document neighbors W_D , represented here by a normalization factor Z . Eq. (5) shows the finally derived k NN-WC model.

$$P(C|D) \approx \frac{1}{Z} f(W_D) \quad (5)$$

4. CONTROVERSY LANGUAGE MODELS

Cramer [5] manually studies patterns of text surrounding specific terms as *controversy*, *dispute*, *scandal*, and *saga* within the Reuters corpus [19], as being indicative of controversy. Motivated by Cramer’s research, we propose a new probabilistic model of controversy that considers how similar

the document’s language is to the one that discusses a range of controversial topics.

Since we are only interested if $P(C|D) > P(NC|D)$ holds, rather than the actual probabilities, we can use rank-safe approximations.

Each $P(C|D)$ and $P(NC|D)$ can be represented using Bayes theorem, which allows us to consider the following odds-ratio:

$$\frac{P(C|D)}{P(NC|D)} = \frac{P(D|C)}{P(D|NC)} \cdot \frac{P(C)}{P(NC)} > 1 \quad (6)$$

Now our test condition can be expressed as:

$$\frac{P(D|C)}{P(D|NC)} > \frac{P(NC)}{P(C)} \quad (7)$$

where for our purposes, we can treat the right hand side as a constant cutoff (since it is independent of the document D), which can be learned with training data. To avoid underflow, we actually calculate the log of this ratio.

$$\log P(D|C) - \log P(D|NC) > \alpha \quad (8)$$

Therefore, we only have to estimate the probabilities $P(D|C)$ and $P(D|NC)$, which we do using the language modeling framework, construction of a language model of controversy L_C , and a non-controversial language model L_{NC} . We make the standard term independence assumption for each word (w) in our document (D), and avoid zero probabilities with linear smoothing. In practice, we estimate both the global language model (L_G) [6] and the non-controversial language model (L_{NC}) as the same by constructing them from the set of all documents.

$$P(D|C) \approx P(D|L_C) = \prod_{w \in D} (\lambda P(w|L_C) + (1 - \lambda)P(w|L_G))$$

$$P(D|NC) \approx P(D|L_{NC}) \approx P(D|L_G) = \prod_{w \in D} P(w|L_G)$$

Here, D_C is a set of controversial documents, and D_{NC} is a set of non-controversial documents, which we estimate in our collections as the background collection, D_{BG} .

$$P(w|L_C) = \frac{\sum_{d \in D_C} tf(w, d)}{\sum_{d \in D_C} |d|}, P(w|L_{NC}) = \frac{\sum_{d \in D_{BG}} tf(w, d)}{\sum_{d \in D_{BG}} |d|}$$

Therefore, to build a language model of controversy, we need to find D_C . We explore WCF features and Cramer-inspired query based models to construct D_C as following:

WCF: Top K Wikipedia articles that have high WCF (e.g, M, C, D) values.

Controversy-indicative terms: Documents that are retrieved by a query believed to indicate controversy. We explore Cramer’s terms [5] as well as manual lexicons from past work [14, 18].

5. EVALUATION

We leverage the dataset introduced in prior work [8] that consists of judgments for 303 webpages¹ from the ClueWeb09 collection². We perform 5-fold cross-validation and report measures on the reconstructed test set.

We implement our probabilistic model based on k NN-WC, the state-of-the-art approach. Rather than use the full text

¹<http://ciir.cs.umass.edu/downloads>

²<http://lemurproject.org/clueweb09/>

Table 1: Wikipedia-Based Controversy Detection Approaches. All LM approaches have significant improvements over their respective k NN-WC counterpart at the $p < 0.05$ level.

Method	WCF	AUC
k NN-WC model §3	M	0.733
k NN-WC model	C	0.743
k NN-WC model	D	0.500†
LM §4	M	0.801
LM	C	0.835
LM	D	0.795

† In the k NN-WC-D approach, no neighbors were found with dispute tags, so it is equivalent to the weak baseline performance of the NO classifier.

of the web pages, we follow the k NN-WC algorithm and use the “TF10” query, where the document is approximated by the ten most frequent terms (excluding the 571-word SMART stopword list).

Prior work reported accuracy; we note that 65% of the 303 documents were non-controversial, so that accuracy does not provide the best view of this dataset. In this work, we primarily present results using the Area Under the Curve (AUC) measure, as we can compare performance without tuning thresholds. Since accuracy was used in prior work, we report it as well. Compared to k NN-WC algorithm, we improve from 0.72 accuracy (as presented in [8]) and 0.737 accuracy (as reproduced) to 0.779 ($p < 0.001$). For our statistical significance tests, we follow in the footsteps of the **PROC**³, and obtain confidence intervals from bootstrap resamples of the predictions.

For each fold, we trained two parameters by grid search: K , the number of top documents to choose, and λ , the smoothing parameter (Table 1). For example, to create our M-score-based language model, we ranked the documents in our Wikipedia collection by their M score, and derived a language model based on the concatenation of the top K documents.

For building Cramer language models, where the relevant document sets were not created by WCF, we used the Galago search engine to rank documents using a query-likelihood retrieval. We explore 6 different corpora as document sources (Table 2). The K highest-scoring documents were then used as our controversial document set: D_C .

6. RESULTS

In Table 1, we present results of our models built around WCF, as introduced in §2.1. All our language modeling approaches are significantly stronger than the k -NN derived approaches. We only report results of WCF features independently because methods of aggregating these features did not improve significantly over the best feature, and these methods were not quite comparable across k NN-WC and LM approaches.

In Table 2, we present an initial exploration of Cramer’s hypothesis that news is able to name and define controversy. While Cramer defined four keywords to be indicative of controversy, we find that “**controversy**” dominates effectiveness on this dataset. We explore these keywords as queries into an expansion corpus, and construct a language model from

³<http://cran.r-project.org/web/packages/PROC>

Table 2: Language Models built from documents relevant to Cramer’s controversial terms [5]. Collection size $|C|$ in millions of documents and type shown for comparison of results. We found that our wiki dataset was significantly better than all others, which had no pairwise differences otherwise.

Expansion Dataset	Type	$ C $	AUC
DBPedia	Wiki	4.6M	0.853
ClueWeb09B (Spam60)	Web	33.8M	0.741
Reuters	News	0.8M	0.745
NYT-LDC	News	1.8M	0.710
Robust04	News	0.5M	0.711
Signal-1M	News	1M	0.710

Table 3: Language Models built from Cramer’s terms and existing lexicons on DBPedia. We find that “controversy” is the most indicative term, and that “saga” is no better than random. Combining terms led to no improvement over “controversy” alone.

Query to build D_C	AUC
controversy	0.856
Roitman [18]	0.823
dispute	0.740
scandal	0.721
Mejova [14]	0.698
saga	0.500

the highest scoring documents for the given query. That language model is then used for classification.

While we were pleasantly surprised by the efficacy of this simple approach, we did not see the best performance in the news corpora [19] used by Cramer, but rather in using DBPedia as the expansion set. We also explored this approach on other news datasets (Robust04, NYT-LDC [20], and Signal1M [4]) but results were statistically equivalent on all news corpora we tried.

Roitman et al. [18] and Mejova et al. [14] present manually-curated lexicons for controversy tasks. We explore their use intrinsically, with Jaccard Similarity between the lexicon and the document terms in Table 4 and as queries to build a language model in Table 3.

Table 4: A comparison of lexicons built manually and through crowd-sourcing in prior work to our automatically derived language models A (*) indicates significant improvement over the best lexicon approach.

Method	Document	AUC
Roitman Lexicon [18]	TF10	0.543
Mejova Lexicon [14]	TF10	0.562
Mejova Lexicon [14]	Full	0.615
Roitman Lexicon [18]	Full	0.695
Cramer Language Model	Full	0.783
WCF Language Model	Full	0.823*
WCF Language Model	TF10	0.835*
Cramer Language Model	TF10	0.856*

7. CONCLUSION

We introduce probabilistic approaches to modeling controversy, by both recasting prior work into a theoretical framework, as well as introducing a new model. Using insights from recent social science research, we motivate and explore the first language modeling approach to detecting controversy. We find that our new approach is statistically

better than prior work, while simultaneously being more efficient. We demonstrate that strongly indicative terms are as helpful for this problem as complicated Wikipedia-based controversy features and more effective than existing lexicons.

8. ACKNOWLEDGMENTS

This work was supported in part by the Center for Intelligent Information Retrieval, in part by NSF grant IIS-0910884, 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 Peter Cramer, Guy Campos and Gonen Dori-Hacohen for thoughtful discussions.

9. REFERENCES

- [1] R. Awadallah, M. Ramanath, and G. Weikum. Harmony and dissonance: Organizing the people’s voices on political controversies. In *WSDM’12*.
- [2] M.-A. Cartwright, E. Aktolga, and J. Dalton. Characterizing the subjectivity of topics. In *SIGIR’09*.
- [3] Y. Choi, Y. Jung, and S.-H. Myaeng. Identifying controversial issues and their sub-topics in news articles. In *Intelligence and Security Informatics*. Springer, 2010.
- [4] D. Corney, D. Albakour, M. Martinez, and S. Moussa. What do a million news articles look like? In *Workshop on Recent Trends in News IR, ECIR’16*.
- [5] P. A. Cramer. *Controversy as news discourse*, volume 19. Springer Science & Business Media, 2011.
- [6] W. B. Croft, D. Metzler, and T. Strohman. *Search engines: Information retrieval in practice*, volume 283. Addison-Wesley Reading, 2010.
- [7] S. Das, A. Lavoie, and M. Magdon-Ismael. Manipulation among the arbiters of collective intelligence: how wikipedia administrators mold public opinion. In *CIKM’13*.
- [8] S. Dori-Hacohen and J. Allan. Automated controversy detection on the web. In *ECIR’15*.
- [9] S. Dori-Hacohen, E. Yom-Tov, and J. Allan. Navigating Controversy as a Complex Search Task. In *Supporting Complex Search Tasks, at ECIR’15*.
- [10] S. L. Gerhart. Do Web search engines suppress controversy? *First Monday*, 9, 2004.
- [11] M. Jang and J. Allan. Improving automated controversy detection on the web. In *SIGIR’16*.
- [12] M. Kacimi and J. Gamper. Diversifying search results of controversial queries. In *CIKM’11*.
- [13] A. Kittur, B. Suh, B. A. Pendleton, and E. H. Chi. He says, she says: conflict and coordination in wikipedia. In *SIGCHI’07*.
- [14] Y. Mejova, A. X. Zhang, N. Diakopoulos, and C. Castillo. Controversy and sentiment in online news. *arXiv preprint arXiv:1409.8152*, 2014.
- [15] E. Pariser. *The Filter Bubble: What the Internet Is Hiding from You*. 2011.
- [16] A.-M. Popescu and M. Pennacchiotti. Detecting controversial events from twitter. In *CIKM’10*.
- [17] H. S. Rad and D. Barbosa. Identifying controversial articles in wikipedia: A comparative study. In *WikiSym 2012*.
- [18] H. Roitman, S. Hummel, E. Rabinovich, B. Sznajder, N. Slonim, and E. Aharoni. On the Retrieval of Wikipedia Articles Containing Claims on Controversial Topics. In *WWW’16*.
- [19] T. Rose, M. Stevenson, and M. Whitehead. The Reuters Corpus Volume 1. In *LREC 2002*.
- [20] E. Sandhaus. The New York Times annotated corpus. *LDC*, 6(12):e26752, 2008.
- [21] T. Yasseri, R. Sumi, A. Rung, A. Kornai, and J. Kertész. Dynamics of conflicts in wikipedia. *PloS one*, 7(6):e38869, 2012.