UvA-DARE (Digital Academic Repository)

Improved and Robust Controversy Detection in General Web Pages Using Semantic Approaches under Large Scale Conditions

Linmans, J.; van de Velde, B.; Kanoulas, E.

DOI

[10.1145/3269206.3269301](https://doi.org/10.1145/3269206.3269301)

Publication date

2018

Document Version

Final published version

Published in

CIKM'18

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Linmans, J., van de Velde, B., & Kanoulas, E. (2018). Improved and Robust Controversy Detection in General Web Pages Using Semantic Approaches under Large Scale Conditions. In *CIKM'18: proceedings of the 2018 ACM International Conference on Information and Knowledge Management : October 22-26, 2018, Torino, Italy* (pp. 1647-1650). The Association for Computing Machinery. <https://doi.org/10.1145/3269206.3269301>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Improved and Robust Controversy Detection in General Web Pages Using Semantic Approaches under Large Scale Conditions

Jasper Linmans
University of Amsterdam
jasperlinmans@gmail.com

Bob van de Velde
University of Amsterdam
rnavandeveld@gmail.com

Evangelos Kanoulas
University of Amsterdam
ekanoulas@gmail.com

ABSTRACT

Detecting controversy in general web pages is a daunting task, but increasingly essential to efficiently moderate discussions and effectively filter problematic content. Unfortunately, controversies occur across many topics and domains, with great changes over time. This paper investigates neural classifiers as a more robust methodology for controversy detection in general web pages. Current models have often cast controversy detection on general web pages as Wikipedia linking, or exact lexical matching tasks. The diverse and changing nature of controversies suggest that semantic approaches are better able to detect controversy. We train neural networks that can capture semantic information from texts using weak signal data. By leveraging the semantic properties of word embeddings we robustly improve on existing controversy detection methods. To evaluate model stability over time and to unseen topics, we assess model performance under varying training conditions to test cross-temporal, cross-topic, cross-domain performance and annotator congruence. In doing so, we demonstrate that weak-signal based neural approaches are closer to human estimates of controversy and are more robust to the inherent variability of controversies.

ACM Reference Format:

Jasper Linmans, Bob van de Velde, and Evangelos Kanoulas. 2018. Improved and Robust Controversy Detection in General Web Pages Using Semantic Approaches under Large Scale Conditions. In *The 27th ACM International Conference on Information and Knowledge Management (CIKM '18), October 22-26, 2018, Torino, Italy*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3269206.3269301>

1 INTRODUCTION & PRIOR WORK

Controversy detection is an increasingly important task. Controversial content can signal the need for moderation on social platforms, either to prevent conflict between users or limit the spread of misinformation. More generally, controversies provide insight into societies [4]. Often, the controversial content is outside the direct control of a platform on which it is shared, mentioned or discussed. This raises the requirement of generally applicable methods to gauge controversial content on the web for moderation purposes. Unfortunately, what is controversial changes, and may lie more in

the way topics are discussed rather than *what* is discussed, making it difficult to detect controversies in a robust fashion. We take the task of controversy detection and evaluate robustness of different methodologies with respect to the varying nature of controversies.

Prior work on detecting controversies has taken three kinds of approaches: 1) **lexical approaches**, which seek to detect controversies through *signal terms*, either through bag-of-word classifiers, lexicons, or lexicon based language models [6]. 2) **explicit modeling** of controversy through platform-specific features, often in Wikipedia or social-media settings. Features such as mutual reverts [14], user-provided flags [2], interaction networks [12] or stance-distributions [8] have been used as platform-specific indicators of controversies. The downside of these approaches is the lack of generalizability due to their platform-specific nature. 3) **matching models** that combine lexical and explicit modelling approaches by looking at lexical similarities between a given text and a set of texts in a domain that provides explicit features [3, 6, 7].

Controversy detection is a difficult task because 1) controversies are *latent*, like ideology, meaning they are often not directly mentioned as controversial in text. 2) Controversies occur across a vast range of topics with varying topic-specific vocabularies. 3) Controversies change over time, with some topics and actors becoming controversial whereas others stop to be so. Previous approaches lack the power to deal with such changes. Matching and explicit approaches are problematic when the source corpus (e.g. Wikipedia) lags after real-world changes [5]. Furthermore, lexical methods trained on common (e.g. fulltext) features are likely to memorize the controversial topics in the training set rather than the 'language of controversy'. Alleviating dependence on platform specific features and reducing sensitivity to an exact lexical representation is paramount to robust controversy detection. To this end, we focus only on fulltext features and suggest to leverage the semantic representations of word embeddings to reduce the vocabulary-gap for unseen topics and exact lexical representations.

The majority of NLP-task related neural architectures rely on word embeddings, popularized by Mikolov et al [11] to represent texts. In essence these embeddings are latent-vector representations that aim to capture the underlying meaning of words. Distances between such latent-vectors are taken to express semantic relatedness, despite having different surface forms. By using embeddings, neural architectures are also able to leverage features learned on other texts (e.g. pretrained word embeddings) and create higher level representations of input (e.g. convolutional feature maps or hidden-states). These properties suggest that neural approaches are better able to generalize to unseen examples that poorly match the training set. We use two often applied network architectures adopting word embeddings, to classify controversy: Recurrent Neural Networks [13] and Convolutional Neural Networks [9] to answer

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 the author(s) 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 '18, October 22-26, 2018, Torino, Italy

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-6014-2/18/10...\$15.00

<https://doi.org/10.1145/3269206.3269301>

the following research question. **RQ:** Can we increase robustness of controversy detection using neural methods?

Currently, there is no open large-size controversy detection dataset that lends itself to test cross-temporal and cross-topic stability. Thus we generate a Wikipedia crawl-based dataset that includes general web pages and is sufficiently large to train and test high capacity models such as neural networks.

2 METHODS

A proven approach in modelling text with neural networks is to use Recurrent Neural Networks (RNNs) which enjoy weight sharing capabilities to model words irrespective of their sequence location. A specific type, the Hierarchical Attention Network (HAN) proposed by [13] makes use of attention to build document representations in a hierarchical manner. It uses bi-directional Gated Recurrent Units (GRUs) [1] to selectively update representations of both words and sentences. This allows the network to both capture the hierarchy from words to sentences to documents and to explicitly weigh all parts of the document relevant during inference.

Recently, Convolutional Neural Networks (CNNs) have enjoyed increasing success in text classification. One such network introduced by [9] looks at patterns in words within a window, such as "Scientology [...] brainwashes people". The occurrences of these patterns are then summarized to their 'strongest' observation (max-pooling) and used for classification. Since pooling is applied after each convolution, the output size of each convolutional operation itself is irrelevant. Therefore, filters of different sizes can be used, each capturing patterns in different sized word windows.

We explore the potential of RNNs and CNNs for controversy detection using both the HAN [13] and the CNN [9] model¹. Similar to [13], each bi-directional GRU cell is set to a dimension of 50, resulting in a word/sentence representation of size 100 after concatenation. The word/sentence attention vectors similarly contain 100 dimensions, all randomly initialized. The word windows defined in the CNN model are set to sizes: 2, 3 and 4 with 128 feature maps each. Each model is trained using mini batches of size 64 and uses both dropout (0.5) and l_2 regularization ($1e-3$) at the dense prediction layer. Both networks use pre-trained embeddings, trained on 100 billion words of a Google News corpus², which are further fine-tuned during training on the controversy dataset. The optimization algorithm used is Adam[10] (learning rate: $1e-3$).

3 EXPERIMENTAL SETUP

3.1 Datasets and evaluation

We use the Clueweb09 derived dataset of [4] for baseline comparison. For cross-temporal, cross-topic and cross-domain training & evaluation, we generate a new dataset based on Wikipedia crawl data³. This dataset is gathered by using Wikipedia's 'List of Controversial articles' overview page of 2018 (time of writing) and 2009 (for comparison with baselines)⁴. Using this as a 'seed' set of controversial articles, we iteratively crawl the 'See also', 'References' and 'External links' hyperlinks up to two hops from the seed list.

The negative seed pages (i.e. non controversial) are gathered by using the random article endpoint⁵. The snowball-sample approach includes general, non-Wikipedia, pages that are referred to from Wikipedia pages. The dataset thus extends beyond just the encyclopedia genre of texts. Labels are assumed to propagate: a page linked from a controversial issue is assumed to be controversial. The resulting dataset statistics are summarized in Table 1.

Table 1: Wikipedia derived dataset statistics. Including the percentages of controversial (i.e. positive labelled) and general (i.e. non-Wikipedia) web pages from the total amount of pages per dataset split.

Set	Seeds	Total	Controversial	General Web
Train	5600	23.703	7.233 (31%)	15.449 (65%)
Validation	200	988	651 (66%)	688 (70%)
Test	200	1.024	654 (64%)	723 (71%)

To be useful as a flagging mechanism for moderation, a controversy detection algorithm should satisfy both Precision and Recall criteria. F1 scores will therefore be used to evaluate this balance. The AUC values are used to measure classification performance in the unbalanced controversy datasets. The test-train split depends on the task investigated and is listed in the results section for the respective task. To test for significant results, all models were evaluated using a bootstrap approach: by drawing 1000 samples with replacements n documents from the test set equal to the test-set size. The resulting confidence intervals based on percentiles provide a measure of significance.

3.2 Baseline models

To compare the results of neural approaches to prior work we implemented the previous state-of-the-art controversy detection method: the language model from [7]. Together with an SVM baseline they act as controversy detection alternatives using only full text features, thus meeting the task-requirements of platform-independence. Note: the implementation of [7] additionally requires ranking methods to select a subset of the training data for each language model. A simplified version of this, excluding the ranking method but using the same dataset and lexicon to select documents as [7], is implemented and included in the baselines comparison section (LM-DBPedia). We also included the same language model trained on the full text Wikipedia pages (LM-wiki). Similarly, for completeness sake, we also include both the state-of-the-art matching model, the TILE-Clique model from [6] and the sentiment analysis baseline (using the state-of-the-art Polyglot library for python⁶) from [3] in the comparison with previous work.

4 RESULTS

4.1 Comparison of results with previous work

Table 2 shows the relative performance of the neural models compared to previous controversy detection methods, evaluated on the Clueweb09 derived dataset of [3] and trained on the Wikipedia data from the same time frame. The TILE-Clique matching model

¹Code available at <https://github.com/JasperLinmans/ControversPy>

²Available at <https://code.google.com/p/word2vec/>

³Script to generate dataset available at: <https://github.com/JasperLinmans/ControversPy>

⁴Available at https://en.wikipedia.org/wiki/Wikipedia:List_of_controversial_issues

⁵Available at <https://en.wikipedia.org/wiki/Special:Random>

⁶<https://github.com/aboSamoor/polyglot>

Model	Precision			Recall			F1			AUC		
	<i>Train/Test: '18/'18</i>	<i>'09/'18</i>	Δ	<i>'18/'18</i>	<i>'09/'18</i>	Δ	<i>'18/'18</i>	<i>'09/'18</i>	Δ	<i>'18/'18</i>	<i>'09/'18</i>	Δ
TfIdf-SVM	0.910	0.941	$\blacktriangle 3\%$	0.689	0.191	$\blacktriangledown 72\%$	0.784	0.317	$\blacktriangledown 60\%$	0.785	0.585	$\blacktriangledown 25\%$
LM	0.651	0.609	$\blacktriangledown 6\%$	0.811	0.550	$\blacktriangledown 32\%$	0.723	0.578	$\blacktriangledown 20\%$	0.600	0.452	$\blacktriangledown 25\%$
CNN	0.930	0.913	$\blacktriangledown 2\%$	0.663	0.564	$\blacktriangledown 15\%$	0.775	0.696	$\blacktriangledown 11\%$	0.888	0.846	$\blacktriangledown 5\%$
HAN	0.871	0.912	$\blacktriangle 5\%$	0.818	0.561	$\blacktriangledown 31\%$	0.844	0.695	$\blacktriangledown 18\%$	0.889	0.845	$\blacktriangledown 5\%$

Table 3: Temporal stability experiment. Results obtained by evaluating on the Wikipedia derived dataset from 2018 by either: models trained on Wikipedia data from 2018 or 2009. Trained on data from the same time frame, the neural models show a slight advantage over the lexical models. Most noticeable however is the drop in performance by the lexical models when trained on older data in terms of Recall and therefore also in terms of F1-score.

outperforms all other models on Precision although this difference is not significant compared to the neural approaches. Similarly, the language model trained on the DBPedia dataset outperforms other models on Recall but shows no significant difference compared to the CNN model. Notably, the neural approaches show comparable results to the TILE-Clique model in terms of F1, demonstrating a balanced performance in terms of Precision and Recall. Furthermore, the CNN model shows a significant improvement compared to the other non neural baselines in terms of the AUC value ($p < 0.05$).

Table 2: Comparison of results with previous work

Model	Precision	Recall	F1	AUC
Sentiment_polyglot	0.448	0.392	0.418	0.612
TfIdf-SVM	0.581	0.208	0.306	0.740
TILE_Clique [6]	0.710	0.720	0.714	0.780
LM-DBPedia [7]	0.415	0.886	0.566	0.730
LM-wiki	0.359	0.808	0.497	0.579
CNN	0.627	0.840	0.718	0.835
HAN	0.632	0.745	0.684	0.823

4.2 Robustness of the model across time

Controversy is expected to change over time. Some issues become controversial, others cease to be so. To investigate robustness of controversy detection models with respect to changes over time, we evaluate model performance in two variants: trained and tested on 2018, or trained on the 2009 Wikipedia data and tested on the 2018 Wikipedia data. Table 3 shows the results for each of the text-based detection models.

Within year, the hierarchical attention model (HAN) outperforms all other models on Recall, F1 and AUC, losing Precision to the CNN and SVM models. However, our main interest is the robustness when a model is trained on a different year (2009) than the test set (2018). These **between** year experiments show a superior score for the HAN model compared to the non-neural models on Recall, and show significant improvements on F1 ($p < 0.05$) and

Figure 1: Average F1 and AUC score of aggregated results for all lexical and neural models.

AUC ($p < 0.05$), losing only to the SVM model on Precision (non significantly). In terms of robustness, we can also take the percentage change between the within year and between year experiment into account (were smaller absolute changes are preferable), shown by the delta values. With regard to temporal sensitivity, the CNN shows the least change across all four metrics. In Figure 1, we show the pooled results for the lexical and neural models to illustrate the overall increase in robustness by neural approaches.

Interestingly, the SVM and HAN model show some unexpected improvement with regard to Precision when applied to unseen timeframes. For both models, this increase in Precision is offset by a greater loss in Recall, which seems to indicate both models ‘memorize’ the controversial topics in a given timeframe instead of the controversial *language*. Overall, the neural approaches seem to compare favorably in terms of cross-temporal stability.

4.3 Robustness of the model across topics

Table 4: Cross-topic stability experiment. Metrics are averaged across 10 leave-on-out topic folds.

Model	Precision	Recall	F1	AUC
TfIdf-SVM	0.793	0.575	0.661	0.829
LM	0.512	0.816	0.629	0.633
CNN	0.840	0.569	0.670	0.842
HAN	0.799	0.716	0.753	0.840

To evaluate robustness towards unseen topics, 10-fold cross validation was used on the top ten largest topics present in the Wikipedia dataset in a leave-one-out fashion. The results are shown in table 4. In line with previous results, the language model scores best on Recall, beating all other models with a significant difference ($p < 0.01$). However in balancing Recall with Precision, the HAN model scores best, significantly outperforming both lexical models in F1 score ($p < 0.05$). Overall, when grouping together all neural and lexical results, the neural methods outperform the lexical models in Precision ($p < 0.01$), F1 ($p < 0.05$) and AUC ($p < 0.01$) with no

significant difference found on the overall Recall scores. These results indicate that neural methods seem better able to generalize to unseen topics.

4.4 Robustness of the model across domains

Most work on controversy has looked into using existing knowledge bases as a source of controversy information [3, 6]. In this paper, we focus on text-based classification methods that do not aim to explicitly link general web pages to their knowledge-base counterparts. Therefore, we are interested in the ability of neural models to generalize beyond their training context. In addition to testing across time and topics, we also investigate robustness to changes in domain. By training only on Wikipedia data, and evaluating only on general web-pages, we look at the ability of the four methods to deal with out-of-domain documents.

The hierarchical attention network shows significantly better results ($p < 0.05$) compared to all other models on F1. Both neural models also outperform both language models on AUC significantly ($p < 0.05$). Precision and Recall are more mixed, with the CNN and SVM outperforming the HAN on Precision and the language model -again- performing best in terms of Recall. Together, the neural methods seem to work best on three out of the four metrics.

Table 5: Cross-domain stability experiment. Metrics are based on models trained on Wikipedia data and tested on general web pages.

Model	Precision	Recall	F1	AUC
TfIdf-SVM	0.718	0.361	0.480	0.638
LM	0.392	0.826	0.532	0.573
CNN	0.743	0.394	0.514	0.755
HAN	0.700	0.604	0.645	0.789

4.5 Human agreement

Lastly, we examine model performance with respect to human annotation using the human annotated dataset of [3]. We assume that models that perform similarly to human annotators are preferable. In Table 6, we present three Spearman correlation metrics to express model congruence with human annotations. Mean annotation expresses the correlation of model error rates with the controversy values attributed to a web page by human annotators, with positive values expressing greater error rates on controversial, and negative values expressing higher error rates on non-controversial pages. Here, the HAN shows most unbiased (closest to zero) performance.

Certainty is the distance of human annotations to the midpoint of the four-point controversy scale, i.e. a score between 0 and 2.5 that expresses how sure annotators are of document (non)controversy. Here, the HAN shows errors most strongly negatively correlated to the certainty of annotators. Finally, annotators disagree on the controversy of some documents, expressed as the standard deviation of their controversy annotations. Again, the HAN model seems preferable, as it's errors are most strongly correlated to annotator disagreement. Overall, the neural methods have less biased performance in relation to (non)controversial documents, correlate more strongly with the certainty of human annotators and are susceptible to errors in similar conditions as when annotators disagree.

Table 6: Spearman's correlations for estimated probability distance from true label. Mean controversy: Average annotator score, certainty: distance from controversy annotation-scale midpoint, disagreement: standard deviation of annotations. Only pages with at least 3 annotations included to ensure sensible agreement metrics, N=128, bolded scores are preferable.

Model	mean annotation	certainty	disagreement
TfIdf-SVM	-0.540	-0.238	0.144
LM-DBPedia	0.633	-0.172	-0.023
CNN	0.348	-0.314	0.138
HAN	0.277	-0.390	0.207

5 CONCLUSION

Controversy detection is a hard task, as it forms a latent concept sensitive to vocabulary gaps between topics and vocabulary shifts over time. We analysed the performance of language model, SVM, CNN and HAN models on different tasks.

First, we have demonstrated that neural methods perform as state-of-the-art tools in controversy detection on the ClueWeb09 [4] based testset, even beating matching models. **Second**, we investigated temporal stability, and demonstrated neural -and especially CNN- robustness in terms of Recall, F1 and AUC performance and stability with train and test sets that are 9 years apart. **Thirdly**, we show that CNN and HAN models outperform the SVM and LM baselines on Precision, F1 and AUC when tested on held-out-topics. **Fourthly**, we show that neural methods are better able to generalize from Wikipedia pages to unseen general web pages in terms of Precision, F1 and AUC. **Lastly**, neural methods seem better in line with human annotators with regard to certainty and disagreement.

REFERENCES

- [1] Cho K. & Bengio Y. Bahdanau, D. 2014. Neural machine translation by jointly learning to align and translate.. In *arXiv preprint arXiv:1409.0473*.
- [2] Sanmay Das, Allen Lavoie, and Malik Magdon-Ismael. 2016. A Manipulation among the arbiters of collective intelligence: How Wikipedia administrators mold public opinion. *ACM Trans. Web* 10, 4, 24:1–24:25.
- [3] Allan J. Dori-Hacohen, S. 2013. Detecting Controversy on The Web.. In *Proceedings of the 22nd ACM CIKM*, 1845–1848.
- [4] Yom-Tov E. Allan J. Dori-Hacohen, S. 2015. Navigating controversy as a complex search task.. In *CEUR Workshop Proceedings*, 1338.
- [5] Odijk D.-& de Rijke M. Graus, D. 2018. The birth of collective memories: Analyzing emerging entities in text streams.. In *JASIST*, 69(9). 773–786.
- [6] Allan J. Jang, M. 2016. Improving Automated Controversy Detection on the Web. In *39th ACM SIGIR Proceedings* 865–868.
- [7] Foley J.-Dori-Hacohen S. Allan J. Jang, M. 2016. Probabilistic Approaches to Controversy Detection.. In *Proceedings of the 25th ACM CIKM*. 2069–2072.
- [8] Myungha Jang, Shiri Dori-Hacohen, and James Allan. 2017. Modeling Controversy within Populations. In *ICTR Proceedings*. ACM, 141–149.
- [9] Y. Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *Proceedings of the 2014 EMNLP*. 1746–1751.
- [10] Jimmy. Kingma, Diederik P. Ba. 2014. Adam: A Method for Stochastic Optimization. In *International Conference for Learning Representations*. arXiv:1412.6980.
- [11] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space.. In *CoRR*, Vol. abs/1301.3781.
- [12] M. Popescu, A.-M. Pennacchiotti. 2010. Detecting controversial events from twitter.. In *Proceedings of the 19th ACM CIKM - CIKM*.
- [13] Dyer D. He X. Smola. A. Hovy. E. Yang. Z., Yang Z. 2016. Hierarchical Attention Networks for Document Classification.. In *Proceedings of NAACL-HLT 2016*. 1480–1489.
- [14] Sumi R. Yasseri, T. 2012. Dynamics of Conflicts in Wikipedia. In *PLOS ONE*, 7(6), 1–12.