

Mitigating Confirmation Bias on Twitter by Recommending Opposing Views

Elisabeth Lex
Graz University of Technology &
Know-Center GmbH, Austria
elisabeth.lex@tugraz.at

Mario Wagner
Graz University of Technology,
Austria
Mario-wagner@gmx.at

Dominik Kowald
Know-Center GmbH,
Austria
dkowald@know-center.at

ABSTRACT

Contemporary politics are often discussed on social media and thus, social media is an important information and news source [2]. Politicians and parties increasingly use social media to spread their ideas, to interact with voters and/or critics and to shape the political debate [6, 7]. However, over the past years, social networks such as Facebook and Twitter have been accused of facilitating the spread of falsehoods, hatred and conspiracy theories [1]. Social networks have also been blamed to create so-called *echo chambers* of like-minded people who are exposed mainly to views that are similar to their existing beliefs [3, 11]. Interestingly, people tend to perceive information so it confirms their existing beliefs, a phenomenon called *confirmation bias* [12]. One countermeasure to confirmation bias is to give more prominence to opposing views [11].

The Present Work. In this work, we propose a content-based recommendation approach to increase exposure to opposing beliefs and opinions. Our aim is to help provide users with more diverse viewpoints on issues, which are discussed in partisan groups from different perspectives. Since due to the *backfire effect* [12], people’s original beliefs tend to strengthen when challenged with counter evidence, we need to expose them to opposing viewpoints at the right time. The preliminary work presented here describes our first step into this direction.

As illustrative showcase, we take the political debate on Twitter around the presidency of Donald Trump. In the 2016 US election, Twitter, among other social networks, played a strong role in the parties’ campaigns and in shaping the political debate. For example, the campaign motto “*Make America Great Again*” of President Trump has been translated in the widely adopted hashtag #MAGA. Quickly, people from all parts of the political spectrum have adopted the hashtag to discuss issues such as e.g., border policies from a variety of viewpoints. Pro-Trump Twitter users have used #MAGA to e.g., express their wish to build a wall on the border to Mexico in combination with #BuildThatWall while contra-Trump users have used #MAGA combined with #NoWall to express their opposing viewpoint on this issue. In other words, people discuss the same issues but from diverse perspectives.

Approach and Dataset. We base our work on Graells et al. [5], who created data portraits of users to determine a user’s stance to the issue *abortion in Chile*. They then used these data portraits to connect people of opposing views. For our work, we crawled a dataset consisting of two partisan groups with two stances: (i) pro-Trump users and (ii) contra-Trump users. To identify those two types of users, we used a set of manually selected hashtags. We selected these hashtags according to their discriminatory power to distinguish between pro and contra-Trump accounts as well as according to their

reach¹. In case of pro-Trump accounts, we used the following hashtags: ‘maga’, ‘tcot’², ‘americafirst’, ‘trumptrain’, ‘presidenttrump’, ‘draintheswamp’, ‘fakenews’, ‘potus’, ‘buildthewall’, ‘presidentelecttrump’. In case of contra-Trump accounts, we used the hashtags: ‘impeachtrump’, ‘theresistance’, ‘nobannowall’, ‘resist’, ‘trumprussia’, ‘impeach45’, ‘nottheenemy’, ‘resistance’, ‘notmypresident’, ‘iamamuslimtoo’, ‘nobannowallnoraid’, ‘fakepresident’, ‘dumprump’, ‘trumplies’.

This procedure resulted in a dataset of 73,868 tweets from 39,698 accounts. Next, we extracted single group accounts, i.e., accounts that strictly use hashtags from only one group – either the pro-Trump or the contra-Trump group since many hashtags appear in both groups. In order to further filter out managed accounts such as news channels, we created boxplots for the number of followers, the number of followees, the number of likes and the number of status updates. All accounts, which were below the first quartile and above the third quartile were filtered out, resulting in 6,913 accounts. We also removed accounts with a non-English user language, resulting in 5,672 accounts. We then downloaded the last 1,000 tweets from these remaining accounts, imported them into Apache Solr³, normalized the texts to lowercase and performed tokenization, as well as stopword removal using the Snowball framework⁴.

Since our initial classification is affected by the multiple meanings of a hashtag (e.g., #MAGA), we followed the method of [5] to reliably classify the users into pro-Trump and contra-Trump users. Thus, we created pro-Trump and contra-Trump issue stance vectors by concatenating all tweets of the users from each stance and extracting trigrams using TF-IDF. This results in our two issue stance vectors. We did the same to create user stances, i.e., we concatenated the tweets of each user u and ran TF-IDF to extract user u ’s trigrams, resulting in a set of 5,672 user vectors. By calculating the cosine similarity between the issue stance vectors and each user vector, we can classify our users into one of the two issue stances. In total, we identified 2,150 pro-Trump users with 2,615,140 pro-Trump tweets and 3,522 contra-Trump users with 3,852,895 contra-Trump tweets.

To determine a user’s personal preference about the stance, we extract the 15 most common trigrams from her tweets, as shown exemplary for a Twitter account in Figure 1.

We then implemented a content-based filtering recommendation approach by exploiting Apache Solr’s MoreLikeThis functionality⁵. For each recommendation, we created a candidate set of 100k tweets, i.e., 50k random pro-Trump tweets and 50k random contra-Trump

¹We assessed them with the service [hashtag.org](http://www.hashtags.org): <http://www.hashtags.org>.

²‘tcot’ stands for ‘top conservatives on Twitter’.

³<https://lucene.apache.org/solr/>

⁴<https://github.com/snowballstem>

⁵We will share the link to our GitHub repository shortly.

Username	#Tweets	Stance	Trigrams (user preferences)
FxgFx	1199	contratrup-stance	'trump pay mexico', 'mexico would pay', 'pay wall meant', 'trump's tax returns', 'showing true face', '#trumpussia #russiagate #resist' 'rep devin nunes', 'health care plan', 'house oversight committee', 'can't wait til', 'get new orders', 'defund planned parenthood', 'make health insurance', 'bring candles back', 'gop members congress'

Figure 1: Example Twitter account “FxgFx”, who has been classified into the contra-Trump stance and her most common trigrams, which serve as proxy for her preferences.

Trigram: “trump pay mexico” for user FxgFx			
Nr	Username	Userstance	Text
1	ckenned	contratrup-stance	RT @SarahLerner: TRUMP: "Mexico will pay for the wall!"
2	CollinDemocrat	contratrup-stance	MEXICO: https://t.co/Sb0Pphp6bA
3	bobedgar208	contratrup-stance	RT @chrislhayes: When he said Mexico was gonna pay for the wall, what he "meant" was homebound senior citizens are gonna pay for the wall.
4	concerned4theUS	Protrump-stance	RT @JoyAnnReid: So "Mexico will pay for the wall" actually meant "we'll cancel healthcare for millions unless YOU pay for the wall."
			Good! Now tell mexico to pay for the wall or else, f**k off! https://t.co/UIxFCwa08

Figure 2: Ex.: Recommendations for Twitter account “FxgFx”.

tweets from our dataset. Then, we computed the cosine similarity between a user’s top trigram and the candidate set. An example recommendation set for the Twitter account “FxgFx” is shown in Figure 2 for her top trigram “trump pay mexico”.

Evaluation. Since we are interested in mitigating confirmation bias by increasing exposure to opposing views, we created four variants of our content-based filtering recommender: (i) “Standard”, which recommends the 10 best matching tweets regardless of the user stance (as shown in Figure 2), (ii) “Pro-Trump”, which only recommends pro-Trump tweets, (iii) “Contra-Trump”, which only recommends contra-Trump tweets, and (iv) “Hybrid”, which recommends half pro-Trump and half contra-Trump tweets (i.e., 5 + 5).

We evaluated these variants with beyond accuracy metrics of recommender systems research, i.e., we computed diversity and serendipity of recommendations [8]. Diversity is measured by the intra-list similarity metric, which sums all pairwise cosine similarities of the items in a given set and calculates the average of the sum. If a set has many similar items, the score is high, if the items are very different, the score is low. Serendipity [4] measures how surprising the recommendations for a user are. In other words, serendipity denotes the distance between recommended items and their expected content.

To further understand how diverse the pro-Trump and contra-Trump users in our dataset are per se, we computed the average topic similarity per user stance. We define the average topic similarity per issue stance as the average pairwise cosine similarity between all users of an issue stance (i.e., pro-Trump and contra-Trump).

Results and Discussion. The results of our evaluation are given in Table 1 calculated for 1,500 randomly chosen pro-Trump users and 1,500 randomly chosen contra-Trump users. For the serendipity metric, as expected, the best results are achieved by the recommendation variant, which recommends tweets from the opposing view. This means that contra-Trump tweets provide the highest serendipity effect for pro-Trump users and vice versa. If both pro-Trump and

Issue stance	Recommendation variant	Serendipity	Diversity
Pro-Trump	Standard	.935	.560
	Pro-Trump	.943	.630
	Contra-Trump	.951	.695
	Hybrid	.946	.728
Contra-Trump	Standard	.924	.441
	Pro-Trump	.957	.728
	Contra-Trump	.925	.487
	Hybrid	.940	.701

Table 1: Our evaluation results with respect to recommendation serendipity and diversity. Here, the “Hybrid” approach recommends half pro-Trump and half contra-Trump tweets.

contra-Trump tweets are mixed, serendipity gets lower. With respect to diversity, in the pro-Trump setting, the best results are achieved by the hybrid variant, which is also the behavior that we expected. However, in the contra-trump setting, we observe a different and rather surprising behavior since the best results here are achieved by the “Pro-Trump” variant of our recommender approach and not the hybrid one. One reason for this could be in the higher average topic similarity of contra-Trump users in our dataset. Interestingly, there is a big difference between both groups since the contra-Trump accounts exhibit an average topic similarity of 44.6% while the pro-Trump accounts exhibit an average topic similarity of only 27.7%. If the user group, in this case contra-Trump, has a high average topic similarity its inherent diversity is lower. Thus, diversity becomes lower if many tweets from a low diversity group are mixed into the recommendations. Consequently, better diversity results could be achieved if fewer (less than 50%) of the more similar contra-Trump tweets and more of the diverse pro-Trump tweets are recommended. How to best engineer this ratio will be part of our future work.

Conclusion and Future Work. We propose a hybrid content-based recommendation approach to mitigate confirmation bias and to help increase exposure to opposing views and beliefs. Our idea is to combine recommendations from partisan groups, i.e., pro-Trump and contra-Trump, into a mixed set that contains tweets from both sides. For evaluation, we turned to beyond-accuracy metrics of recommender systems, i.e., diversity and serendipity and we found that our approach lets us boost both metrics. In the longer run, we think of the problem at hand as an optimization task to achieve a trade-off between recommendation accuracy, serendipity and diversity.

For future work, we will research on the optimal mixing strategy to create the hybrid recommendations from pro-Trump and contra-Trump users. We also will study communication patterns in our dataset to better understand partisan and cross-partisan interactions. Also, we will research when to best add cross-partisan recommendations to increase the chance of recommendation acceptance [3]. One idea is to detect shifts in attention and user focus e.g., with the SUSTAIN [9] or ACT-R algorithm [10], or shifts in conversational strength or sentiment. Finally, we also plan to verify our findings in larger Twitter data samples and with different political topics.

Acknowledgments. This work is funded by the Know-Center GmbH (Austrian COMET program) and the H2020 project AFEL (grant agreement: 687916).

Keywords. Confirmation Bias; Tweet Recommendations; Diversity; Serendipity; Polarization; Hybrid Recommendations

REFERENCES

- [1] Jon Berkeley. 2017. Do social media threaten democracy? (November 2017). <https://www.economist.com/leaders/2017/11/04/do-social-media-threaten-democracy> [Online; posted 04-November-2017].
- [2] Michael Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Goncalves, Filippo Menczer, and Alessandro Flammini. 2011. Political Polarization on Twitter. (2011). <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2847>
- [3] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2017. Reducing Controversy by Connecting Opposing Views. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining (WSDM '17)*. ACM, New York, NY, USA, 81–90. DOI : <http://dx.doi.org/10.1145/3018661.3018703>
- [4] Mouzhi Ge, Carla Delgado-battenfeld, and Dietmar Jannach. 2010. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In *In RecSys 10*. 257.
- [5] E. Graells-Garrido, M. Lalmas, and D. Quercia. 2013. Data Portraits: Connecting People of Opposing Views. *ArXiv e-prints* (Nov. 2013). arXiv:cs.HC/1311.4658
- [6] Sounman Hong and Sun Hyoung Kim. 2016. Political polarization on twitter: Implications for the use of social media in digital governments. *Government Information Quarterly* 33, 4 (2016), 777 – 782. DOI : <http://dx.doi.org/https://doi.org/10.1016/j.giq.2016.04.007>
- [7] Andreas Jungherr. 2016. Twitter use in election campaigns: A systematic literature review. *Journal of information technology & politics* 13, 1 (2016), 72–91.
- [8] Marius Kaminskis and Derek Bridge. 2017. Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 7, 1 (2017), 2.
- [9] Simone Kopeinik, Dominik Kowald, Ilire Hasani-Mavriqi, and Elisabeth Lex. 2017. Improving Collaborative Filtering Using a Cognitive Model of Human Category Learning. *The Journal of Web Science* 2, 4 (2017), 45–61. DOI : <http://dx.doi.org/10.1561/106.00000007>
- [10] Dominik Kowald, Subhash Chandra Pujari, and Elisabeth Lex. 2017. Temporal Effects on Hashtag Reuse in Twitter: A Cognitive-Inspired Hashtag Recommendation Approach. In *Proceedings of the 26th International Conference on World Wide Web (WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 1401–1410. DOI : <http://dx.doi.org/10.1145/3038912.3052605>
- [11] Azi Lev-On and Bernard Manin. 2009. Happy Accidents: Deliberation and Online Exposure to Opposing Views. *Online Deliberation: Design, Research and Practice*, Todd Davies, Seeta Gangadharan, eds., Forthcoming. (October 2009). Available at SSRN: <https://ssrn.com/abstract=1481869>.
- [12] Craig Silverman. 2017. The Backfire Effect. (June 2017). https://archives.cjr.org/behind_the_news/the_backfire_effect.php [Online; posted 17-June-2011].