Purple Feed: Identifying High Consensus News Posts on Social Media

Mahmoudreza Babaei[†], **Juhi Kulshrestha**[†], **Abhijnan Chakraborty**[†‡]
Fabrício Benevenuto[†*], **Krishna P. Gummadi**[†], **Adrian Weller**[**]
[†]Max Planck Institute for Software Systems (MPI-SWS), Germany
[‡]Indian Institute of Technology Kharagpur, India
[*]Universidade Federal de Minas Gerais, Brazil
[**]University of Cambridge & Alan Turing Institute, United Kingdom

Abstract

Although diverse news stories are actively posted on social media, readers often focus on news which reinforces their pre-existing views, leading to 'filter bubble' effects. To combat this, some recent systems expose and nudge readers toward stories with different points of view. One example is the Wall Street Journal's 'Blue Feed, Red Feed' system, which presents posts from biased publishers on each side of a topic. However, these systems have had limited success.

In this work, we present a complementary approach which identifies high consensus 'purple' posts that generate similar reactions from both 'blue' and 'red' readers. We define and operationalize *consensus* for news posts on Twitter in the context of US politics. We identify several high consensus posts and discuss their empirical properties. We present a highly scalable method for automatically identifying high and low consensus news posts on Twitter, by utilizing a novel category of *audience leaning based features*, which we show are well suited to this task. Finally, we build and publicly deploy our 'Purple Feed' system (twitter-app.mpi-sws.org/purple-feed), which highlights high consensus posts from publishers on both sides of the political spectrum.

Introduction

A growing number of people rely on social media platforms, such as Twitter and Facebook, for their news and information needs (Lichterman 2010; Teevan, Ramage, and Morris 2011). As a result, the users themselves play a role in selecting the sources from which they consume information, overthrowing traditional journalistic gatekeeping (Shoemaker, Vos, and Reese 2009). To cope with the huge amount of information available, most social media platforms, together with third-party developers, have deployed information selection and retrieval systems which help users discover and read interesting news and information.

Within our societies, there are many topics for which different subgroups hold opposing ideological positions. For example, there are primarily two distinct political affiliations in the US: Republicans (the 'red' group) and Democrats (the 'blue' group). Social media platforms provide a wide variety of news sources covering this ideological spectrum, yet many users largely limit themselves to news stories which reinforce their pre-existing views. This *selective exposure*, where red users read red news and blue users read blue news, may lead to a more politically fragmented, less cohesive society (Liu and Weber 2014). Further, this effect is often amplified by social media platforms which recognize users' preferences and thence recommend more red news to red users and more blue news to blue users. While this approach may work well for recommending consumer goods such as movies or music, there are concerns that such stilted news selections limit exposure to differing perspectives and can lead to the formation of 'filter bubbles' or 'echo chambers' (Bakshy, Messing, and Adamic 2015; Bozdag 2013; Flaxman, Goel, and Rao 2016; Pariser 2011), resulting in a worrying increase in *societal polarization* (Sunstein 2002; Schkade, Sunstein, and Hastie 2007).

To counter this polarization, a number of systems intended to promote diversity have been proposed. These systems deliberately expose users to different points of view by showing red news to blue users, and blue news to red users; or by showing both red and blue news to both red and blue user groups. The hope is to nudge users to read viewpoints which disagree with their own (Munson, Lee, and Resnick 2013; Park et al. 2009). A prominent example is the Wall Street Journal's 'Blue feed, Red feed' system (Keegan 2017), which presents posts from the most extreme news publishers on Facebook, with the aim of showing diametrically opposed perspectives on news stories.

Unfortunately, however, such systems have had limited success. While some diversity-seeking users enjoy the added perspectives, many users either ignore or reject disagreeable points of view (Munson and Resnick 2010). Indeed, by confronting users with the most radical posts from the other ideological side, such systems may even *increase* polarization by encouraging users to retreat to a more entrenched version of their initial position (Lord and Ross 1979; Miller et al. 1993; Munro and Ditto 1997).

In this work, we propose a complementary approach by identifying and highlighting news posts which are likely to evoke similar reactions from the readers, irrespective of their political leanings. We define these 'purple' news posts to be those with *high consensus*, *i.e.,* having a general agreement in their readers' reactions to them. We propose that these high consensus purple stories could be recommended to both red and blue users, evoking a more unified response across society, which we hope might lead to lower segregation in information consumption (Chakraborty et al. 2017), and might help to promote greater understanding and cohesion among people. In Table 1, we show a sample of red, blue and purple news stories about the dismissal of FBI director James Comey by President Trump, to highlight the differences between the three types of stories.[1]

Given this context, we investigate the following questions:

1. How can we define the consensus of news posts in order to operationalize the identification of high consensus purple posts?
2. Do helpful purple news posts exist on social media?
3. How do purple posts compare with low consensus (blue or red only) posts?
4. Can we automate the identification of consensus of news posts on social media in order to discover purple posts?

Contributions

1. We begin by defining and operationalizing the concept of consensus of news posts in terms of general agreement in readers' reactions.
2. We use human judgments to generate a ground truth dataset of high and low consensus news posts on social media, and observe that a substantial amount of high consensus purple posts are posted by news publishers on social media (perhaps surprisingly, even by politically extreme publishers).
3. We analyze the properties of high and low consensus posts empirically, finding that both types of tweets are equally popular with users (*i.e.,* garner similar no. of retweets) and also cover similar topics. Further, we observe that high consensus purple posts tend to provide more cross-cutting exposure to views than low consensus posts.
4. To identify high consensus purple news posts automatically, we propose a novel class of features of social media posts on Twitter, which we term *audience leaning based features*. These features describe the distribution of the political leanings of audience subgroups interacting with a post – namely the retweeters and repliers of a post. Intuitively, most of the retweeters are likely to be supportive of it, while repliers have a higher likelihood of opposing it. Additionally, the followers of the publisher of the post also form a passive audience subgroup for the post. We use these audience leanings as features to capture the degree of consensus that a social media post is likely to have. We show that our proposed features are well suited to help identify high and low consensus tweets automatically with high accuracy, leading to significantly better performance than can be achieved using previously proposed publisher based and content based features.

Our work provides a fresh tool which we hope may help to burst filter bubbles, encourage healthier interaction between population subgroups, and lead to a more cohesive society.

[1]See en.wikipedia.org/wiki/Dismissal_of_James_Comey.

Red Posts	Fox News: @seanhannity: "The real reason that President #Trump fired James Comey is because the former @FBI Director was incompetent.". https://t.co/YlP6SLOfgw
	Fox News: @POTUS: "All of the Democrats, I mean, they hated Jim Comey. They didn't like him, they wanted him fired".https://t.co/1ebOtqfIOc
Blue Posts	Salon: Comey firing coverage shows right-wing media has lost it's grip on reality https://t.co/DC6cAYEDoX
	CNN: Pres. Trump's firing of FBI Director James Comey is a "grotesque abuse of power," legal analyst Jeffrey Toobin says http://cnn.it/2q1FQd4
Purple Posts	NYTimes: He was fired by President Trump. Where does James Comey go next? https://t.co/loXwc5aNFd
	Politico: What happens to Comey's investigative files that have already been gathered? A former FBI special agent weighs in.politi.co/2qbBJfj

Table 1: Sample "red", "blue" and "purple" news posts about the event of FBI director James Comey's dismissal by President Trump.

Consensus Definition and Measurement

A key step of our work is to understand whether news posts with high consensus exist in social media. To verify this, first we need to operationalize the concept of 'consensus' for news posts, i.e., to provide a definition for consensus that allows one to measure it, both empirically and quantitatively. Second, we need to construct ground truth datasets to measure consensus of real news posts in social media. Next, we describe how we performed these steps.

Operationalizing consensus for news posts

According to the Oxford English Dictionary, consensus is defined as "a general agreement".[2] Inspired by this definition, we consider a post to have *high consensus if there is a general agreement in readers' reaction to it, irrespective of their own political leaning*. Specifically, in the context of US politics, a post would have high consensus if the reaction of Democrat readers to the post is similar to the reaction of Republican readers. For a given social media post, we measure the reaction of Democrats and Republicans as whether the readers agree or disagree with the content of a post. Formally, we measure the amount of consensus as

$$\text{consensus} = 1 - \left|\frac{\#D_{disagree}}{\#D} - \frac{\#R_{disagree}}{\#R}\right| \quad (1)$$

where $\#D_{disagree}$ and $\#R_{disagree}$ respectively denote the number of Democrats and Republicans who disagree with the post, while $\#D$ and $\#R$ are the total number of Democrats and Republicans.[3] A consensus value closer to 1 indicates that both Democrats and Republicans disagreed with it to similar extents, thereby indicating high consensus; while a value closer to 0 is indicative of low consensus.

Note that there is no unique way to measure consensus. In addition to Equation 1, it can also be measured in terms of *attitude polarization indices* such as coherence, divergence, intensity, and parity (Persily 2015). The common requirement for these indices is *attitude response data*, which in

[2]See en.oxforddictionaries.com/definition/consensus.

[3]Considering the fraction of readers from each side who disagree with a post implicitly takes into account the fraction of readers who have neutral or favorable reactions to it.

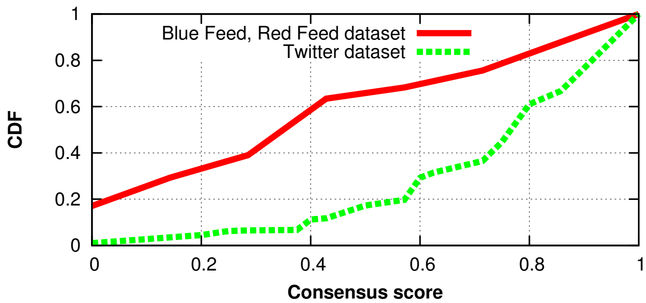

Figure 1: Distribution (CDF) of consensus values of news posts from the two datasets.

our context is provided by the support and negative response from the readers. We use our definition presented in Equation 1 due to its simplicity, while still effectively capturing the nuances of consensus as in the other measures.

Measuring consensus of news posts on social media Using our definition of consensus, we conducted an Amazon Mechanical Turk (AMT) experiment to quantify the consensus of two distinct datasets of news posts:
(i) *Blue Feed, Red Feed dataset* - Using the Wall Street Journal's 'Blue Feed, Red Feed' system (Keegan 2017), we collected the top 10 posts from each of the liberal and conservative sides for the queries "Trump" and "healthcare", giving us a total of 40 news posts.
(ii) *Twitter dataset* - We also collected 40 news posts tweeted by each of the following 10 news publishers with well known political biases varying from liberal to neutral to conservative: Slate, Salon, New York Times, CNN, AP, Reuters, Politico, Fox News, Drudge Report, and Breitbart News; giving us a total of 400 posts. These news posts were collected during the one week period of 9th to 15th May, 2017.

In this experiment, we only recruited AMT workers from the US and at the end of the experiment, we also collected their political leanings. We showed every news post to workers and asked them for their reaction to the post by selecting one out of three options – agreement, neutral or disagreement.[4] After the experiment, applying Equation 1 to responses from an equal number (seven) of Democrat and Republican AMT workers, we computed the consensus values for the news posts in our two datasets.

Figure 1 shows the distribution (CDF) of consensus values for the news posts from our datasets. We observe that news posts from the 'Blue Feed, Red Feed' dataset are skewed towards lower values of consensus, indicating that the readers from the two different parties have different reactions to them; whereas the random news posts from the 10 publishers from the Twitter dataset have a noticeable skew towards higher consensus. Our observations suggest that while news outlets on social media do publish posts with varying degrees of consensus, systems such as 'Blue Feed, Red Feed', which highlight posts from extremely biased news outlets, tend to pick lower consensus content which leads readers with different leanings to react differently. In this work, we make a case to promote high consensus posts in order to increase exposure to ideologically cross-cutting content, which may help in lowering societal polarization.

[4] We presented the questions in the following form:
Tweet: "Salon: Can anyone find an economist who thinks Trumps tax cuts will pay for themselves? https://t.co/Hz6JvQHvXB"
Question: Do you agree with this tweet?
(a) Agree, (b) Neither agree nor disagree, or (c) Disagree.

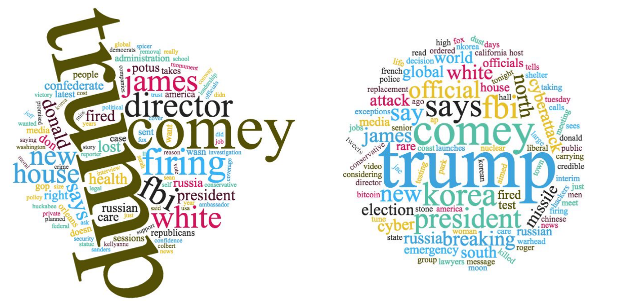

(A) Low consensus tweets (B) High consensus tweets

Figure 2: Topical coverage of low & high consensus tweets.

Empirical Study of Consensus of News Posts

Given that news posts with high consensus do exist, next we conduct an empirical study on consensus of news posts on social media. Our main goal in this section is to understand if news posts with high consensus are interesting to users (*i.e.,* do they become popular), if they cover a wide range of topics, and if they expose users to relatively more cross-cutting content. We refer to the 100 news posts with the highest consensus values in our ground truth Twitter dataset as *high consensus* news posts, and the 100 tweets with lowest consensus values as *low consensus* news posts.

To what extent do high and low consensus news posts become popular? To measure popularity, we count the number of retweets for high and low consensus tweets. On average, high consensus tweets are retweeted 158 times, whereas low consensus tweets are retweeted 177 times, i.e. slightly more often but the numbers are very close. We observe a similar pattern when we compare the median number of retweets of high consensus tweets (93) with low consensus tweets (89), indicating that both high and low consensus tweets engage their readers to similar extents. This suggests that recommendation systems which highlight high consensus tweets would feature content which is of similar popularity to that generated by systems which highlight low consensus tweets.

Do high and low consensus news posts cover different topics? To verify whether high and low consensus tweets cover similar (or very different) topics, we present the 100 most common words for both sets of tweets in Figure 2. From the figure, it is evident that although both sets do cover popular political topics (*e.g.* 'Trump', 'Comey', 'FBI' and other topics associated with FBI director James Comey's dismissal), high consensus tweets are topically more diverse and also contain posts on non-US centric political topics (*e.g.* 'North Korea') and other more niche topics (*e.g.* 'jobs', 'cyberattack').

Feature Category	Features
Publisher based	Number of followers/friends/tweets Average number of retweets/replies/favorites Political leaning, Language, Location
Tweet based	Bag of words, Creation time Number of retweets/replies/favorites

Table 2: Features used in prior work. The three most important features from each category are highlighted in blue.

@CNN You mean, like the UNFOUNDED claims of Russian collusion? You people are typically selective in your bias pro? https://t.co/CESkVpIZOk
@nytimes His actions were disgraceful. Being fired does not make him a sympathetic figure. He affected the outcome? https://t.co/bIbiuj2CJJ
@BreitbartNews I just wonder, what motivates these libtards... https://t.co/mzpBIKdPr4
@CNN hey fakenews do some homework, get out of office! Every illegal that get a drivers license is registered to vote dem! I'd card, regs!
@AP Jews are so desperate to take over Syria that they will make up anything.

Table 3: Random sample of replies for tweets in our dataset.

Do high consensus posts lead to more exposure to ideologically cross-cutting content? To investigate whether highlighting high consensus tweets leads to higher exposure to ideologically cross-cutting contents, we examine whether the higher consensus tweets have relatively more retweets from the users of opposite leaning (with respect to the publisher's leaning), when compared to lower consensus tweets. This analysis is motivated by the reasoning that as users of opposite leaning retweet the publisher's tweets, more opposite leaning users from these users' neighborhoods would get exposed to them, leading to higher exposure to cross-cutting content for users, and potentially lower polarized news consumption on social media.

To validate whether our reasoning holds, we consider a particular tweet to have *high cross-cutting exposure* if the number of opposite leaning retweeters for this tweet is higher than the baseline number of opposite leaning retweeters of its publisher (computed as the average across 100 random tweets of the publisher). When we rank the tweets by their consensus values and compare the top and bottom 10% tweets, we find that a much larger fraction (45%) of high consensus tweets have high cross-cutting exposure than low consensus tweets (30%), indicating that high consensus tweets indeed lead to higher exposure to cross-cutting content.

Identifying High and Low Consensus News Posts on Social Media

After empirically exploring the consensus of social media news posts, we now turn our attention towards *automatically identifying* high and low consensus news posts, which can scale up to cover a large number of news publishers on Twitter. In this section, we first briefly discuss different features of social media posts that have been applied in prior prediction and classification tasks. Then, we propose and validate a novel class of *audience leaning based features* which are

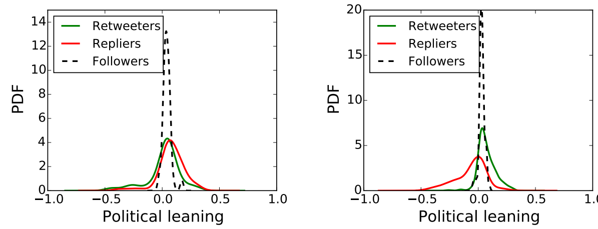

(A) High consensus tweet (B) Low consensus tweet

Figure 3: Distributions of political leanings of different audiences for the following news posts: (A) High consensus: "Trump ordered emergency meeting after global cyber attack: official http://reut.rs/2r6Qkt8" posted by Reuters, (B) Low consensus: "Michelle Obama criticizes Trump administration's school lunch policy http://cnn.it/2qckHwZ" posted by CNN .

Category	Features
Followers	# Dem/Rep/Neu, Sum/Avg/Median/Skew of PL Sum(PL) of Dem/Rep/Neu, PLD
Retweeters	# Dem/Rep/Neu, Sum/Avg/Med/Skew of PL Sum(PL) of Dem/Rep/Neu, PLD of baseline Avg #Dem/Rep/Neu in baseline, PLD χ^2 Distance bw PLD[Retweeters] of tweet & baseline
Repliers	# Dem/Rep/Neu, Sum/Avg/Med/Skew of PL Sum(PL) of Dem/Rep/Neu, PLD of baseline Avg #Dem/Rep/Neu in baseline, PLD χ^2 Distance bw PLD[Repliers] of tweet & baseline
Combination	χ^2 Distance bw PLD[Repliers] and PLD[Retweeters] χ^2 Distance bw PLD[Repliers] and PLD[Followers] χ^2 Distance bw PLD[Retweeters] and PLD[Followers]

Table 4: Audience leaning based features. In the table, Dem, Rep, and Neu denote Democrat, Republican, and Neutral respectively, PL denotes political leaning, and PLD denotes the distribution of political leanings. Baselines are computed by taking average of PLD across all tweets. Most important features are highlighted in blue.

ideally suited for our consensus identification task.

Features used in prior work

Prior works on classification and prediction tasks for social media posts have mostly used two broad types of features: *publisher-based*, and *tweet-based* features. For instance, the political leaning of the publisher has been used to quantify the tweet's leaning (Kulshrestha et al. 2017), or the leaning of news story URLs being shared by them on Facebook (Bakshy, Messing, and Adamic 2015). Others have used tweet-based features for predicting the relevance of a tweet for a topic (Tao et al. 2012), to rank tweets (Duan et al. 2010), or to quantify to what extent a tweet is interesting (Naveed et al. 2011). Many other studies have combined both publisher- and tweet-based features for various tasks including predicting future retweets (Petrovic, Osborne, and Lavrenko 2011; Suh et al. 2010), and even predicting users' personality traits (Golbeck, Robles, and Turner 2011). Table 2 shows the features from each class which we are aware were used previously.

Our proposed audience leaning based features

We propose a novel class of *audience leaning based* features, which to our knowledge have not previously been used

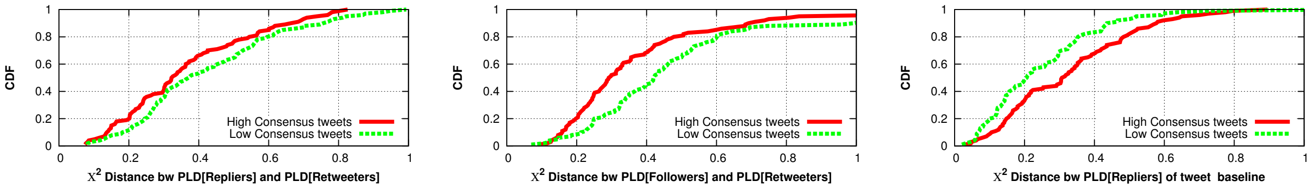

(A) χ^2 Distance between PLD[Retweeters] & PLD[Repliers]

(B) χ^2 Distance between PLD[Retweeters] & PLD[Followers]

(C) χ^2 Distance between PLD[Repliers] & publisher baseline PLD[Repliers]

Figure 4: Distributions of χ^2 distance between different audience political leaning distributions for 25% tweets with highest and lowest consensus values.

for predicting and classifying tweet properties. We use these features to identify high and low consensus posts on Twitter. For every tweet, there are three types of audiences:
(i) *Followers* of the publisher of the tweet – they are the passive supporters of the post (on average 67% of followers are of the same political leaning as the publisher[5]),
(ii) *Retweeters* of the tweet – they are more active supporters of the post (on average 78% of retweeters are of the same leaning as the publisher), and
(iii) *Repliers* to the tweet – they are usually a mix of users supporting or opposing the news post (on average 35% of repliers are of the opposite leaning to the publisher). In Table 3, we show a random sample of replies from our Twitter dataset, and notice that many of them oppose either the news content or the publisher.

We hypothesize that we can use the political leaning distributions of the three audiences of a post to quantify whether different readers of a post are having similar reactions to it (*i.e.,* to measure consensus). To demonstrate our hypothesis, we select one high consensus and one low consensus post for which we computed consensus values using AMT workers' judgments, and then computed the political leaning distributions of the three audiences.

Inferring political leaning of Twitter users is a research challenge on its own, and beyond the scope of this work. We adopt the methodology proposed in (Kulshrestha et al. 2017), which returns the political leaning of a Twitter user in the range $[-1.0, 1.0]$, with scores in $[-1.0, -0.03)$ indicating Republican leaning, $[-0.03, 0.03]$ indicating neutral and $(0.03, 1.0]$ indicating Democrat leaning. In Figure 3, we plot the political leaning distributions of the three audiences, for a high consensus and a low consensus post.

We can observe that there is a striking difference between the audience leaning distributions of high and low consensus tweets in Figure 3. For the high consensus tweet, these distributions are much more similar than for the low consensus tweet. More interestingly, retweeters typically being supporters, have similar political leaning distribution as the followers of the publishers (for both types of posts). However, for a lower consensus post, repliers being opposers, have a different distribution. Therefore, we find that the *degree of similarity of the leaning distributions of the audiences* of the post contains a useful signal to approximate the consensus for a post (*i.e.,* the similarity in reaction of readers of different leanings). We compute the χ^2 distances between the leaning distributions of the different audiences to capture their similarities. In Figure 4, we show the distribution of these χ^2 values for high and low tweets. The difference in the distributions for the high and low consensus posts give evidence for the discriminative power of these features. Building upon these observations, we construct a number of audience leaning based features by utilizing the political leanings of the three types of audiences of a tweet. Table 4 lists all such features, which we use in this work.

[5]Followers of famous politicians (e.g., President Trump) indeed include many users (such as journalists) from both ends of the political spectrum, who may not necessarily support him or his views.

Experimental Evaluation

We first describe our experimental setup, then present our results for the aforementioned categories of features.

Experimental setup: We use supervised learning approaches to identify whether a news tweet has high consensus or low consensus using the features described in the previous section. For setting up the classifiers, we first need a ground truth dataset of high and low consensus tweets. We use the consensus values computed using AMT workers' judgments for the Twitter dataset described previously and label the top 25% consensus value tweets as high consensus, and bottom 25% tweets as low consensus tweets. We use this set of 200 labeled tweets as our ground truth dataset.

Using the features described earlier, we apply four different types of supervised learning classifiers for our task of tweet consensus classification: Linear SVM, Naive Bayes, Logistic Regression and Random Forest classifiers. While using textual features of the tweets, we follow a two step approach as described in (Chakraborty et al. 2016; 2018):
(i) first, we treat the textual features as bag-of-words and use Naive Bayes classifier to predict the class using these textual features, and (ii) then we input these prediction outputs of Naive Bayes classifier as features (along with our other features) to the different classifiers as the second step.

For training our classifiers, we use 5-fold cross-validation. In each test, the original sample is partitioned into 5 sub-samples, out of which 4 are used as training data, and the remaining one is used for testing the classifier. The process is then repeated 5 times, with each of the 5 sub-samples used exactly once as the test data, thus producing 5 results. The entire 5-fold cross validation was then repeated 20 times with different seeds used to shuffle the original dataset, thus producing 100 different results. The results reported are averages of the 100 runs, along with the 90% confidence in-

Classifier	Different feature categories				
	Publisher based (P)	Tweet based (T)	P and T	Audience leaning based (A)	P, T, and A
Logistic Regression	0.58 ± 0.008	0.58 ± 0.008	0.68 ± 0.009	0.72 ± 0.012	0.72 ± 0.011
Linear SVM	0.58 ± 0.008	0.58 ± 0.008	0.68 ± 0.009	0.72 ± 0.012	0.72 ± 0.011
Naive Bayes	0.59 ± 0.007	0.57 ± 0.015	0.60 ± 0.01	0.66 ± 0.015	0.66 ± 0.012
Random Forest	0.58 ± 0.008	0.57 ± 0.01	0.64 ± 0.01	0.67 ± 0.015	0.67 ± 0.017

Table 5: Average accuracies and 90% confidence intervals for different categories of features used for predicting consensus of news tweets. Our proposed audience leaning based features perform best for this news post consensus classification task.

terval. Also, we use feature ranking with recursive feature elimination that prunes out the insignificant features by obtaining their importance from the supervised techniques.[6]

Experimental results: We successively implemented the different classifiers first using features from each category separately, and then by combining the features from different categories. Accuracies are shown in Table 5. We observe that the tweet-based features have the worst performance. This poor performance is most likely due to the short size of the tweets, which often means that there is very little information in the tweet text and it is hard to understand them without also inspecting the content of weblink, photograph or video included in the tweet. The performance of publisher-based features is better than that of tweet-based features. The political leaning of the publisher is found to be the most important feature for this category, and while it helps, it does not perfectly capture the notion of consensus. When we combine publisher- and tweet-based features, there is an improvement in performance.

Next, we examine the performance of our proposed audience leaning based features and find it to perform the best amongst the three categories of features. Digging deeper, we find that we correctly classified 74% of high consensus tweets and 70% of low consensus tweets. We find χ^2 distance between the repliers' and retweeters' leaning distribution to be the most important feature, matching the intuition we built earlier in the paper. In fact, even when we combine the three categories of features, we do not find a performance gain over using the audience leaning based features alone. This is because when we inspect the 10 most important features out of all the categories, the top 7 most important features (highlighted in Table 4 in blue) are from our proposed category of audience leaning based features, highlighting how well suited they are for our consensus identification task.

Conclusion

To minimize the possibility of social media users getting trapped in 'echo chambers' or 'filter bubbles', prior works have proposed to introduce diversity in the news that users are consuming (Munson, Lee, and Resnick 2013; Park et al. 2009; Keegan 2017). Often, such approaches which highlight the most belief challenging news, increase the chances of users rejecting them, thereby defeating the original purpose (Munson and Resnick 2010; Lord and Ross 1979; Miller et al. 1993; Munro and Ditto 1997). In this paper, we propose a complementary approach to inject diversity in users' news consumption by highlighting news posts which evoke similar reactions from different readers, irrespective of their own political leanings.

Towards that end, to our knowledge, we made the first attempt to define and operationalize consensus of news posts on social media. Subsequently, we compared several properties of high and low consensus news posts and found them to be equally popular, and covering similar topics. Additionally, we observed that high consensus posts lead to higher cross-cutting exposure for the users. Next, utilizing our proposed novel class of audience leaning based features, we developed a method to automatically infer the consensus of news posts on Twitter.

Finally, using our proposed consensus inference method, we publicly deployed "Purple Feed" – a system which highlights high consensus posts from different news outlets on Twitter. With "Purple Feed", the users can view the high consensus tweets posted by both Republican-leaning and Democrat-leaning media outlets during the last one week.[7] Users can also view both high and low consensus posts posted by individual publishers.[8]

In future, we plan to conduct a large scale characterization study of news posts and publishers on social media, and to evaluate the impact of showing high consensus news posts on the users. We believe that our work on identifying high consensus news posts could be integrated with different information retrieval mechanisms on social media, and could be useful for designing mechanisms for mitigating filter bubble and echo chambers, for reducing fragmentation in news consumption, and for encouraging healthy debate on diverse issues on social media platforms.

Acknowledgments: We thank the anonymous reviewers and the conference participants whose suggestions helped to improve the paper. A. Chakraborty is a recipient of Google India PhD Fellowship and Prime Minister's Fellowship Scheme for Doctoral Research, a public-private partnership between Science & Engineering Research Board (SERB), Department of Science & Technology, Government of India and Confederation of Indian Industry (CII). A. Weller acknowledges support from the David MacKay Newton research fellowship at Darwin College, The Alan Turing Institute under EPSRC grant EP/N510129/1 & TU/B/000074, and the Leverhulme Trust via the CFI. F. Benevenuto acknowledges Fapemig, CNPq, Capes, and Alexander von Humboldt Foundation.

[6] See http://scikit-learn.org/stable/modules/generated/sklearn.feature_selection.RFE.html.

[7] Available at http://twitter-app.mpi-sws.org/purple-feed/.

[8] For instance, high and low consensus tweets posted by New York Times can be viewed at: http://twitter-app.mpi-sws.org/purple-feed/app-tweet-1.php?query=NYTimes.

References

Bakshy, E.; Messing, S.; and Adamic, L. A. 2015. Exposure to ideologically diverse news and opinion on facebook. *Science* 348(6239):1130–1132.

Bozdag, E. 2013. Bias in algorithmic filtering and personalization. *Ethics and Information Technology* 15(3):209–227.

Chakraborty, A.; Paranjape, B.; Kakarla, S.; and Ganguly, N. 2016. Stop clickbait: Detecting and preventing clickbaits in online news media. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE.

Chakraborty, A.; Ali, M.; Ghosh, S.; Ganguly, N.; and Gummadi, K. P. 2017. On quantifying knowledge segregation in society. *arXiv preprint arXiv:1708.00670*.

Chakraborty, A.; Luqman, M.; Satapathy, S.; and Ganguly, N. 2018. Editorial algorithms: Optimizing recency, relevance and diversity for automated news curation. In *The 2018 Web Conference Companion*, WWW'18. ACM.

Duan, Y.; Jiang, L.; Qin, T.; Zhou, M.; and Shum, H.-Y. 2010. An empirical study on learning to rank of tweets. In *Proceedings of the 23rd International Conference on Computational Linguistics*, COLING '10, 295–303. ACL.

Flaxman, S.; Goel, S.; and Rao, J. M. 2016. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly* 80(S1):298–320.

Golbeck, J.; Robles, C.; and Turner, K. 2011. Predicting personality with social media. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*, CHI EA'11, 253–262. ACM.

Keegan, J. 2017. Blue feed, red feed - see liberal facebook and conservative facebook, side by side. http://graphics.wsj.com/blue-feed-red-feed/.

Kulshrestha, J.; Eslami, M.; Messias, J.; Zafar, M. B.; Ghosh, S.; Gummadi, K. P.; and Karahalios, K. 2017. Quantifying search bias: Investigating sources of bias for political searches in social media. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, CSCW '17, 417–432. ACM.

Lichterman, J. 2010. New Pew data: More Americans are getting news on Facebook and Twitter. *http://tinyurl.com/News-on-Social-Media*.

Liu, Z., and Weber, I. 2014. Is twitter a public sphere for online conflicts? a cross-ideological and cross-hierarchical look. In *International Conference on Social Informatics*, 336–347. Springer.

Lord, C. G., and Ross, L. 1979. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology* 2098–2109.

Miller, A. G.; McHoskey, J. W.; Bane, C. M.; and Dowd, T. G. 1993. The attitude polarization phenomenon: Role of response measure, attitude extremity, and behavioral consequences of reported attitude change. *Journal of Personality and Social Psychology* 64(4):561.

Munro, G. D., and Ditto, P. H. 1997. Biased assimilation, attitude polarization, and affect in reactions to stereotype-relevant scientific information. *Personality and Social Psychology Bulletin* 23(6):636–653.

Munson, S. A., and Resnick, P. 2010. Presenting diverse political opinions: How and how much. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, 1457–1466. ACM.

Munson, S.; Lee, S.; and Resnick, P. 2013. Encouraging reading of diverse political viewpoints with a browser widget. In *Proceedings of the 7th International Conference on Weblogs and Social Media, ICWSM 2013*, 419–428. AAAI.

Naveed, N.; Gottron, T.; Kunegis, J.; and Alhadi, A. C. 2011. Bad news travel fast: A content-based analysis of interestingness on twitter. In *Proceedings of the 3rd International Web Science Conference*, WebSci '11, 8:1–8:7. ACM.

Pariser, E. 2011. *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Group.

Park, S.; Kang, S.; Chung, S.; and Song, J. 2009. Newscube: delivering multiple aspects of news to mitigate media bias. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 443–452. ACM.

Persily, N. 2015. *Solutions to political polarization in America*. Cambridge University Press.

Petrovic, S.; Osborne, M.; and Lavrenko, V. 2011. RT to win! predicting message propagation in twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, ICWSM'11. AAAI.

Schkade, D.; Sunstein, C. R.; and Hastie, R. 2007. What happened on deliberation day? *California Law Review* 95(3):915–940.

Shoemaker, P. J.; Vos, T. P.; and Reese, S. D. 2009. Journalists as gatekeepers. *Handbook of journalism studies* 73.

Suh, B.; Hong, L.; Pirolli, P.; and Chi, E. H. 2010. Want to be retweeted? large scale analytics on factors impacting retweet in twitter network. In *Proceedings of the 2010 IEEE Second International Conference on Social Computing*, SOCIALCOM '10, 177–184. IEEE Computer Society.

Sunstein, C. R. 2002. The law of group polarization. *Journal of Political Philosophy* 10(2):175–195.

Tao, K.; Abel, F.; Hauff, C.; and Houben, G.-J. 2012. What makes a tweet relevant for a topic? In *#MSM- Workshop on Making Sense of Microposts*, volume 838 of *CEUR Workshop Proceedings*, 49–56.

Teevan, J.; Ramage, D.; and Morris, M. R. 2011. #twittersearch: a comparison of microblog search and web search. In *WSDM '11 Proceedings of the fourth ACM international conference on Web search and data mining*, 35–44. ACM.