

Toward a Better Understanding of Emotional Dynamics on Facebook

Fabiana Zollo¹(✉), Borut Sluban², Igor Mozetič³, and Walter Quattrociocchi¹

¹ Ca' Foscari University of Venice, Venice, Italy
{fabiana.zollo,w.quattrociocchi}@unive.it

² University of Zurich, Zurich, Switzerland
borut.sluban@uzh.ch

³ Jožef Stefan Institute, Ljubljana, Slovenia
igor.mozetic@ijs.si

Abstract. On online social media users tend to aggregate in *echo chambers*, where they shape and reinforce their worldview by discussing and interacting with like-minded people. Such a scenario fosters misinformation spreading, which may influence public opinion. To determine the main factors behind narratives' emergence, characterizing polarization dynamics and users' emotional response to social contents is, thus, crucial. In this paper, we address such a challenge by looking at two different and contrasting narratives, science and conspiracy. We introduce a new metric, the *bipolarity*, and show how it can help in finding non-trivial proxies of the debate's polarization. Our approach may provide interesting insights for a better understanding of both emotional and polarization dynamics on online social media.

1 Introduction

In the latest years, several concerns have been posed about the outcomes that online debates may have in shaping public opinion and impacting real-world processes. Previous studies showed that on online social media users tend to select and interpret information that is already coherent to their system of beliefs and that the main driver of contents diffusion is confirmation bias [1–4]. Users aggregate into the so-called echo chambers, i.e., groups of like-minded people, where they reinforce and polarize their opinions and ignore dissenting information [5, 6]. Such a scenario dramatically reduces the potential benefits coming from the exposure to different points of view and fosters misinformation spreading. Also, discussion with like-minded people has been proved to let individuals take on a more extreme position in the end [7]. The phenomenon is alarming, to the extent that since 2013, the World Economic Forum (WEF) has been placing the global danger of massive digital misinformation at the core of other technological and geopolitical risks, ranging from terrorism, to cyber attacks, up to the failure of global governance [8]. When people are misinformed, they hold wrong beliefs neglecting factual evidence, and public opinion may be influenced negatively.

To (further) complicate matters, it has been shown that people usually tend to resist facts, holding inaccurate, factual beliefs confidently [9]. Also, corrections frequently fail to reduce misperceptions [10] and act as a *backfire effect*. Misinformation spreading motivated both researchers [11, 12] and major corporations, such as Google and Facebook [13], to analyze the problem and possibly provide solutions. Users' polarization has been shown to be a crucial aspect of online debates [14–16]. Moreover, Facebook users tend to confine their attention on a limited set of pages (*selective exposure*), thus determining a sharp community structure among news outlets [17]. Similar patterns have been observed both around the Brexit [18] and the Italian Constitutional Referendum [19] debates. Therefore, a better understanding of polarization dynamics and users' emotional response to social contents may be crucial to determine the main factors behind narratives' emergence on online social media. In this paper we address emotional and polarization dynamics inside and across two different and contrasting narratives, science and conspiracy. Following the approach used in [1, 5, 6], we collected all conspiracy and scientific news sources of the US Facebook over a time span of five years (2010–2014). Notice that we do not focus on the quality of information. The main difference between the two categories of pages, indeed, consists in the possibility of verifying their content. While it is relatively easy to verify scientific information (e.g., authors of the study, institutions/universities where it took place, peer-review processes), conspiracy-like content is more difficult to verify, being usually built upon suspect information. The list of pages has been built with the support of very active debunking groups (see Sect. 2 for further details)¹. We downloaded all public posts (with the related likes and comments) of 83 scientific and 330 conspiracy pages. In addition, we identified 66 Facebook pages aiming at debunking conspiracy theories. Then, we applied sentiment analysis techniques to the comments of such posts, and study their aggregated sentiment. The sentiment classification is based on a supervised machine learning method. For this purpose, a subset of all comments was first manually annotated for sentiment and then used as a training set for a learning algorithm to build a domain-specific sentiment classification model. The model is then applied to associate each comment with a sentiment score, which is intended to express the emotional attitude of Facebook users when posting comments. Although many studies focused on sentiment analysis techniques to social media [20–22], in this work we address the emotional dynamics around different and contrasting narratives, with a special focus on misinformation spreading. Following a previous approach [5], we extend our analysis by introducing new metrics and studying a specific test case involving users' emotional reaction to fact-checking. In particular, the main contribution of this paper is two-fold: (1) We introduce a new metric, the *bipolarity*, to better characterize emotional dynamics of online debates; (2) We show how such a metric can help in finding non-trivial proxies of the debate's polarization. Our approach may provide interesting insights for a better understanding of both emotional and polarization dynamics around different narratives on online social media.

¹ For the complete list please refer to [6].

2 Methods

2.1 Data Collection

We use an approach based on [1, 6, 23] to identify two main categories of pages: conspiracy news – i.e. pages promoting contents *neglected* by mainstream media – and science news. To define the space of our investigation we also received help from Facebook groups very active in debunking conspiracy theses. Pages were categorized according to their content and self-description into three main categories:

1. *Conspiracy*: pages disseminating controversial information, most often lacking supporting evidence and sometimes contradictory of the official news (i.e., conspiracy theories)
2. *Science*: pages aiming at disseminate scientific information, often including scientific institutions and scientific press, and having the main mission of diffusing scientific knowledge
3. *Debunking*: pages active in debunking false rumors online

Notice that we focus on the possibility of verifying information, rather than on its truth value. For example, conspiracy pages often claim that their mission consists in informing people about topics neglected by mainstream media deliberately. Also, we use debunking pages as testbed for the efficacy of fact-checking efforts. To our knowledge, the final dataset is the complete set of all scientific, conspiracy, and debunking information sources active on the US Facebook. For each page, we collected all the posts over a time span of five years (from January 2010 to December 2014). Data collection was performed exclusively by means of the Facebook Graph API [24], which is publicly available and accessible through one’s personal Facebook account. We collected only data available on public Facebook pages that are public entities. Users’ content contributing to such entities is also public unless users’ privacy restrictions specify otherwise. The exact breakdown of the data is shown in Table 1. Notice that we restrict our analysis only to the posts of our dataset having at least a comment, a like, and a share.

2.2 Sentiment Classification

We assess the opinions, emotions or attitudes of Facebook users towards different posts by sentiment classification of their corresponding comments. The sentiment classification is automatic, by applying a supervised machine learning method. For this purpose, a subset of all comments was first manually annotated for sentiment and then used as a training set for a learning algorithm to construct a domain specific sentiment classification model. We properly evaluate the performance of the classifier, and also compare it to a classifier constructed from a large set of annotated tweets, and to a lexicon-based sentiment classifier.

Table 1. Breakdown of the Facebook dataset. Number of pages, posts, likes, comments, likers, and commenters for science, conspiracy, and debunking pages.

	Total	Science	Conspiracy	Debunking
Pages	479	83	330	66
Posts	682,455	262,815	369,420	50,220
Posts ^a	319,150	129,029	171,706	18,415
Likes	613,515,345	463,966,540	145,388,131	4,160,674
Comments	30,889,614	22,093,692	8,307,643	488,279
Comments ^a	24,907,891	17,672,115	6,854,091	381,685
Likers	52,753,883	40,466,440	19,386,132	744,023
Commenters	9,812,332	7,223,473	3,166,725	139,168
Commenters ^a	8,320,395	6,033,019	2,755,281	120,269

^aPosts with at least a like, a comment, and a share, and related comments and commenters

2.2.1 Data Annotation

It has been shown that the quality of sentiment classification models depends much more on the quality and size of training data than on the type of the model trained [25]. Hence, rather than using a generic sentiment classification model it was essential to obtain sentiment annotated training data that captures the specifics of our studied domain, and use it to construct a domain specific sentiment model.

We sampled almost 25,000 comments from *Conspiracy* and *Science* pages and engaged five human annotators to provide sentiment labels for individual comments. The annotators were instructed to label each comment by a *negative*, *neutral*, or *positive* label. The guideline given was to estimate the emotional attitude of the user when posting a comment to Facebook. During the sentiment labeling process, a fraction of the comments (about 20%) was intentionally annotated twice, either by the same annotator or by two different annotators. Multiple annotations of a comment by the same annotator are used to compute the self-agreement, and multiple annotations by different annotators to compute the inter-annotator agreement (abbreviated as inter-agreement). Self- and inter-agreement of annotators are considered as proxies for the annotation quality/consistency. Inter-agreement also provides an upper bound of the performance achievable by the constructed classification model [25]. In Table 2 we summarize the results of manual sentiment annotations. Self- and inter-annotator agreements are measured in terms of Krippendorff's *Alpha* reliability [26]. We also show comparative results on a related domain, namely generic English Twitter posts, from which training data can be used for sentiment model construction [25]. The number of annotated tweets is considerably larger (set size), but the quality of annotations is considerably lower (*Alpha*). For comparison, the last line gives the annotation results of our previous Italian Facebook study [5].

Table 2. The self- and inter-annotator agreements in terms of Krippendorff’s α . Twitter(EN) denotes a body of generic English tweets, labeled with sentiment. Facebook(IT) gives a comparable results for the Italian Facebook study.

Dataset	Annotations		Self-agreement		Inter-agreement	
	Set size	Annotators	Set size	α	Set size	α
Twitter(EN)	87,428	9	3,392	0.739	12,214	0.613
Facebook(US)	24,312	5	2,546	0.916	4,009	0.810
Facebook(IT)	19,642	25	640	0.854	3,262	0.673

2.2.2 Classification Models

The 24,312 comments manually annotated with sentiment labels were used as training data for the construction of a sentiment classification model. We employ a Support Vector Machine (SVM) [27] supervised machine learning algorithm to build the model. In particular, we implemented an extended version of the SVM algorithm called *TwoPlaneSVMbin* [25], which assumes the ordering of sentiment classes and implements ordinal classification. It combines two SVM models: one to distinguish between the *negative-or-neutral* and *positive* comments, and another to distinguish between the *negative* and *neutral-or-positive* comments. Furthermore, it partitions the space around both hyperplanes into bins, and computes the distribution of the training examples in individual bins. During the classification of all the comments from our Facebook dataset, the distances from both hyperplanes determine the appropriate bin, but the class is determined by the majority class in the bin.

Another sentiment classifier, constructed from a large number of labeled English Twitter posts (abbreviated Twitter(EN), Table 2, [25]) was used for comparison. The tweets are shorter than typical Facebook comments, but the training set is larger. There are also lexicon-based approaches to sentiment classification. In this case, the sentiment of a text is computed from the set of sentiment-bearing words identified in the text. For example, authors [28,29] collected human sentiment assessment of 10,000 common words, each labeled 50 times. We used their English sentiment lexicon and method to implement another, lexicon-based sentiment classifier. The performance of the three constructed classifiers is evaluated on the same training set by 10-fold cross validation. We apply three performance measures: *classification accuracy*, F_1 score averaged over the negative and positive classes, and Krippendorff’s α . α takes into account ordering of sentiment values and classification by chance. When a model perfectly classifies the data, $\alpha = 1$, and when the model is no better than chance classification, $\alpha = 0$. For a detailed description of the selected measures, deemed appropriate for the evaluation of the three-valued sentiment classifiers, please refer to [25].

Table 3 shows the classification performance of the three sentiment classifiers. In terms of all three performance measures, the domain-specific Facebook(US) classifier outperforms the other two. The Twitter(EN) classifier was constructed

by the same SVM algorithm, but from a larger training set where the annotation quality is lower. On the other hand, the lexicon-based classifier is appropriate for much longer texts, where a considerable number of sentiment-bearing words occurs. Note that the “training” set for the lexicon-based classifier is the number of most frequent unique, sentiment-bearing words, that were manually annotated. These results confirm the importance of constructing domain- and even discourse-specific classifiers if one aims at high quality sentiment classification. The last line in Table 3 gives comparable results for our previous Italian Facebook study [5]. In the current study we paid special attention to the higher quality of annotations, which is reflected in better classifier performance. However, the performance (in terms of *Alpha*) is still considerably lower than the inter-annotator agreement (compare to Table 2). In our experience, the number of annotated posts should approach 100,000 when most of the vocabulary is covered and the classifiers can reach the level of performance of human annotators. This is costly in terms of human resources needed, so one usually settles for lower performance at lower costs.

Table 3. Evaluation results of the sentiment classifiers.

Sentiment classifier	Training set size	Performance measures		
		<i>Accuracy</i>	$\overline{F_1}$	<i>Alpha</i>
Lexicon-based	10,222	0.503 ± 0.009	0.543 ± 0.015	0.402 ± 0.021
Twitter(EN)	87,428	0.623 ± 0.018	0.613 ± 0.022	0.528 ± 0.030
Facebook(US)	24,312	0.654 ± 0.015	0.685 ± 0.013	0.589 ± 0.021
Facebook(IT)	19,642	0.648 ± 0.013	0.655 ± 0.012	0.562 ± 0.022

3 Results and Discussion

3.1 Emotional Dynamics Inside Echo Chambers

Previous studies [3, 5] show that users tend to select information that is coherent to their preexisting judgments –i.e., by confirmation bias– and to form polarized groups of like-minded people –i.e., echo chambers– which negatively influence users’ emotions. In this work, we start our analysis by looking at the emotional behavior of users inside the two echo chambers, science and conspiracy. Through the sentiment analysis and classification task we are able to associate each comment of our dataset to a sentiment score $\in [-1, 1]$, where -1 is negative, 0 is neutral, and 1 is positive. Notice that we restrict our analysis to the posts of our dataset having at least a comment, a like, and a share (see Sect. 2 for further details about methods and data collection). In Fig. 1 we plot the Probability Density Function (PDF) of the sentiment score for all the comments of science and conspiracy pages. We observe that distributions for both categories are similar and present three main peaks by values -1 , 0 , and 1 .




Fig. 1. Probability Density Function (PDF) of the sentiment score for all comments on both categories.

Then, we investigate if the distribution of the sentiment changes when focusing directly on posts and users. We define the sentiment score of a post (respectively, user) as the average sentiment score of all its (respectively, her/his) comments. Figure 2 shows the PDFs of the mean sentiment score of posts (left) and users (right) both for science and conspiracy echo chambers. For posts, the distribution moves towards neutral values, suggesting that comments of each post are almost equally divided into positive and negative. However, this could also depend on a high presence of neutral comments. When looking at the mean sentiment score of users, instead, the distributions turn out to be pretty similar to those of comments, showing three pronounced peaks at values -1 , 0 , and 1 .

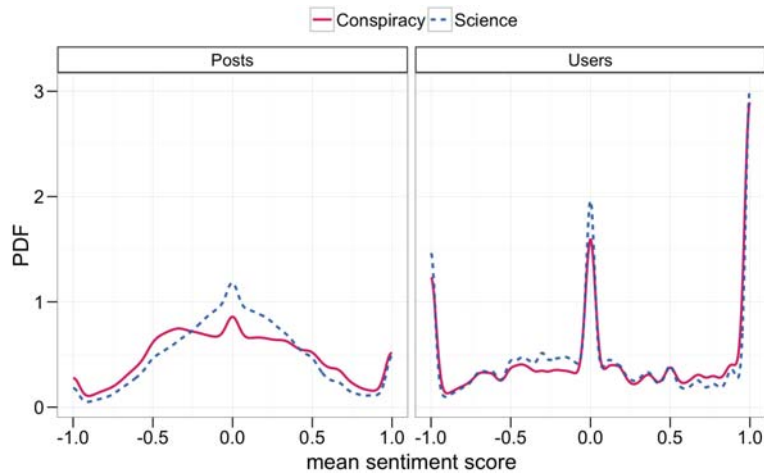


Fig. 2. Probability Density Function (PDF) of mean sentiment score of all posts (left) and users (right) on science and conspiracy pages.

To understand if users' sentiment changes with respect to their engagement in one of the two communities, we focus on the sentiment of polarized users.

Precisely, we define a user to be polarized towards science (respectively, conspiracy) if she has left more than 95% of her likes on science (respectively, conspiracy) posts. Figure 3 shows the PDFs of the mean sentiment score of polarized users for both categories. We observe that both distributions show the same peaks as before, although users polarized towards conspiracy tend to be more negative than those polarized towards science.

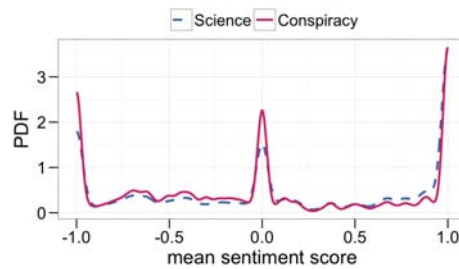


Fig. 3. Probability Density Function (PDF) of mean sentiment score of polarized users.

Nevertheless, such a preliminary analysis is not sufficient to understand the emotional dynamics of the debate. We need to go more in depth and look for new measures that could represent a good proxy of the debate's polarization.

3.2 Emotional Polarization

To further investigate the emotional dynamics around opposing narratives, we use two measures, *subjectivity* and *bipolarity*, which may provide interesting cues to quantify the *emotional polarization* of the debate. In particular, our aim is to understand if such measures can help to determine both if a specific narrative can elicit different (and maybe opposite) reactions in its users and users' emotional attitude within the social network.

Let N denote the total number of comments of the post (respectively, user) and neg , neu , pos the number of negative, neutral, and positive comments, respectively ($N = neg + neu + pos$). The *sentiment score* and *subjectivity* are defined [30] as follows:

$$score = \frac{pos - neg}{N} \quad subjectivity = \frac{pos + neg}{N},$$

where *subjectivity* = 1 if the post (or user) has no neutral comments; vice versa, *subjectivity* = 0 when all the comments are neutral.

Figure 4 shows the PDFs of the subjectivity of posts and polarized users for both categories. We may notice that there is a remarkable difference between posts and users: the first show a visible trend towards 1, denoting a low presence of neutral comments, while the latter present two accentuated peaks at values 0 and 1, showing that most part of users' comments are either neutral or

just negative and/or positive. In the latter case, thus, users' tend to be particularly passionate in their engagement with the social network, showing opposing sentiments in their comments, and a general attitude to avoid neutrality.

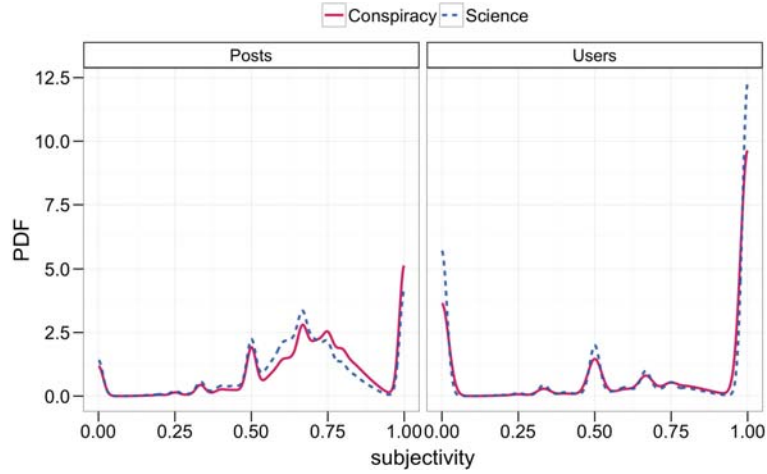


Fig. 4. Probability Density Function (PDF) of subjectivity of posts (left) and polarized users (right).

However, subjectivity alone is not sufficient to capture the complexity of emotional dynamics around the two narratives. To strengthen our analysis, in this paper we introduce a new metric, the *bipolarity*, defined as the product of sentiment *score* and *subjectivity*:

$$bipolarity = score * subjectivity = \frac{pos^2 - neg^2}{N^2},$$

where *bipolarity* = 0 when *pos* = *neg* or *neu* = *N*, and *bipolarity* → 1 (respectively, -1) when *pos* ≫ *neg* (respectively, *neg* ≫ *pos*) and *neu* ≪ *N*.

Figure 5 shows the PDFs of the bipolarity of posts and polarized users for both categories. We find that posts bipolarity shows a trend towards 0 values for both communities, more pronounced for science, although two minor peaks are noticeable at values -1 and 1. Such an observation is particularly meaningful if linked to posts subjectivity depicted in Fig. 4: indeed, since posts subjectivity tends to 1, we know that the presence of neutral comments tends to be very low; consequently, the fact that posts bipolarity goes towards 0 is due to an almost perfect balance of both positive and negative comments. Thus, from an emotional point of view, we can infer that the debate around posts is highly polarized – i.e., elicits different and opposing sentiments among its commenters – on both science and conspiracy pages. Users' bipolarity, instead, shows three pronounced peaks at values -1, 0, and 1. Taking also into consideration the distribution

reported in Fig. 4, such results suggest again that either all comments of a users are neutral, or they are positive and negative, with a clear prevalence of one of the extreme classes, suggesting a strong polarization in users' emotional behavior on the social network.

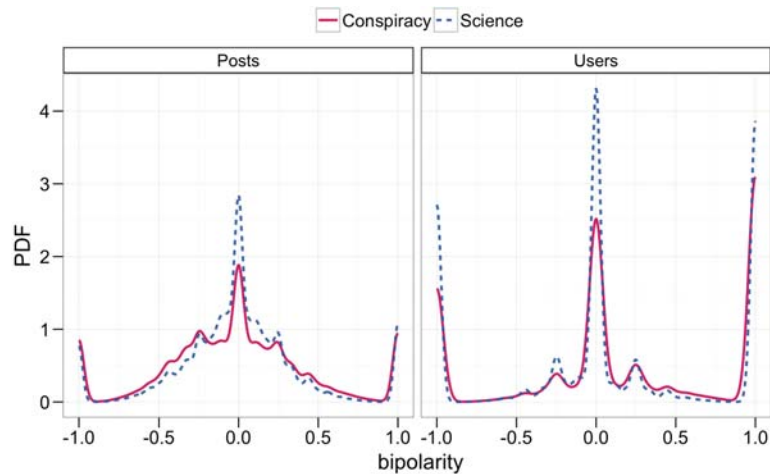


Fig. 5. Probability Density Function (PDF) of bipolarity of posts (left) and polarized users (right).

3.3 Case Study: Emotional Response to Fact-Checking

Previous work [6] showed that fact-checking seem to be ineffective against rumor spreading. Researchers analyzed users' interaction with 50K Facebook debunking posts, showing that (1) such posts remain mainly confined to the scientific echo chamber, (2) only few conspiracy users engage with corrections and their liking and commenting rates on conspiracy posts increases after the interaction, and (3) the sentiment expressed by users when commenting is mainly negative. Now, our aim is to use our new metrics—e.g., subjectivity and bipolarity—to better understand the emotional dynamics induced by fact-checking attempts. Figure 6 shows the PDFs of users' mean sentiment score (left), subjectivity (center), and bipolarity (right), by category. When considering the mean sentiment score, we may notice that science users and others show very similar patterns. More interestingly, the sentiment tends towards negative values, and not only for conspiracy users. However, subjectivity and bipolarity can tell us more. Indeed, distributions of values for subjectivity move towards one, denoting a very low presence of neutral comments. Also, when looking at the bipolarity, we find two high peaks by values -1 and 0 : either lots of comments are negative or there is an equal presence of positive and negative comments. Thus, the debate around

fact-checking appears to be *emotionally* polarized. Such a dominant negativity, together with the emotional polarization, provide interesting clues, suggesting that the way in which debunkers communicate may not be the most appropriate, and that people tend to hold and defend their beliefs and opinions, even when factual evidence is provided.

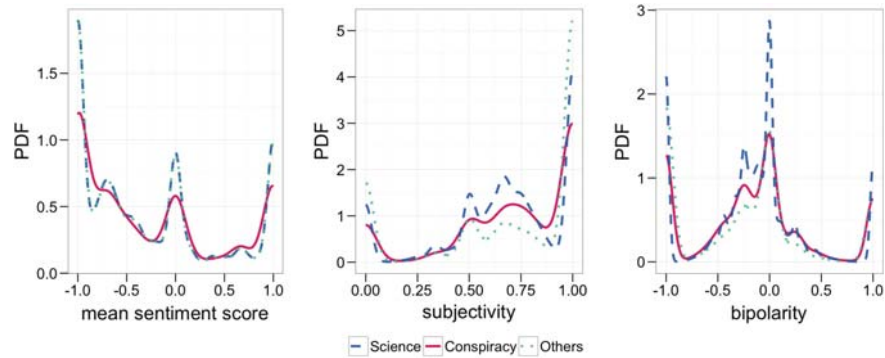


Fig. 6. Probability Density Functions (PDFs) of users’ mean sentiment score, subjectivity, and bipolarity on debunking posts.

4 Conclusions

In this work we applied sentiment analysis techniques to address the emotional dynamics around different and contrasting narratives. We collected data from all conspiracy and scientific pages on Facebook over a time span of five years. We also identified debunking pages aiming at contrasting fake news spreading on the platform. By means of sentiment analysis techniques, we were able to associate each comment of our dataset with a sentiment score and, thus, to analyze the emotional attitude of Facebook users inside both the conspiracy and the scientific echo chamber. Thanks to the introduction of a new metric, the bipolarity, and by its combination with sentiment score and subjectivity, we were able to capture meaningful aspects of users’ behavior in online debates. Our results may provide interesting insights for a better understanding of emotional dynamics around different narratives on online social media.

Acknowledgements. The authors acknowledge financial support from H2020 FET project DOLFINS (grant no. 640772), and the Slovenian Research Agency (research core funding no. P2-103). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

References

1. Bessi, A., Coletto, M., Davidescu, G.A., Scala, A., Caldarelli, G., Quattrociocchi, W.: Science vs Conspiracy: Collective narratives in the age of (mis)information. *PLoS ONE*, **10**(2) (2015)
2. Bessi, A., Petroni, F., Del Vicario, M., Zollo, F., Anagnostopoulos, A., Scala, A., Caldarelli, G., Quattrociocchi, W.: Homophily and polarization in the age of misinformation. *Eur. Phys. J. Spec. Topics* **225**(10) (2016)
3. Del Vicario, M., Bessi, A., Zollo, F., Fabio Petroni, F., Scala, A., Caldarelli, G., Stanley, H.E., Quattrociocchi, W.: The spreading of misinformation online. *PNAS* **113**(3) (2016)
4. Del Vicario, M., Vivaldo, G., Bessi, A., Zollo, F., Scala, A., Caldarelli, G., Quattrociocchi, W.: Echo chambers: emotional contagion and group polarization on Facebook. *Sci. Rep.* **6** (2016)
5. Zollo, F., Kralj Novak, P., Del Vicario, M., Bessi, A., Mozetič, I., Scala, A., Caldarelli, G., Quattrociocchi, W.: Emotional dynamics in the age of misinformation. *PLoS ONE* **10**(9) (2015)
6. Zollo, F., Bessi, A., Del Vicario, M., Scala, A., Caldarelli, G., Shekhtman, L.M., Havlin, S., Quattrociocchi, W.: Debunking in a world of tribes. *PLoS ONE* **12**(7) (2017)
7. Sunstein, C.R.: The law of group polarization. *J. Political Philosophy* **10**(2) (2002)
8. Howell, W.L.: Digital wildfires in a hyperconnected world. WEF Global Risks Report (2013)
9. Kuklinski, J.H., Quirk, P.J., Jerit, J., Schwieder, D., Rich, R.F.: Misinformation and the currency of democratic citizenship. *J. Politics* **62**(3) (2000)
10. Nyhan, Brendan, Reifler, Jason: When corrections fail: the persistence of political misperceptions. *Polit. Behav.* **32**(2), 303–330 (2010)
11. Backstrom, L., Kleinberg, J., Lee, L., Danescu-Niculescu-Mizil, C.: Characterizing and curating conversation threads: Expansion, focus, volume, re-entry. In: Proceedings of WSDM (2013)
12. Romero, D.M., Meeder, B., Kleinberg, J.: Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on Twitter. In: Proceedings of the 20th International Conference on World Wide Web (2011)
13. First draft coalition (2017). <https://firstdraftnews.com>
14. Adamic, L.A., Glance, N.: The political blogosphere and the 2004 US election: divided they blog. In: Proceedings of the 3rd International Workshop On Link Discovery. ACM (2005)
15. Bessi, A., Zollo, F., Del Vicario, M., Ngelo Puliga, M., Scala, A., Caldarelli, G., Brian U., Quattrociocchi, W.: Users polarization on Facebook and Youtube. *PLoS ONE* **11**(8) (2016)
16. Ugander, J., Backstrom, L., Marlow, C., Kleinberg, J.: Structural diversity in social contagion. *PNAS* **109**(16) (2012)
17. Schmidt, A.L., Zollo, F., Del Vicario, M., Bessi, A., Scala, A., Caldarelli, G., Stanley, H.E., Quattrociocchi, W.: Anatomy of news consumption on Facebook. *PNAS* (2017)
18. Del Vicario, M., Zollo, F., Caldarelli, G., Scala, A., Quattrociocchi, W.: Mapping social dynamics on Facebook: the brexit debate. *Soc. Netw.* **50** (2017)

19. Del Vicario, M., Gaito, S., Quattrociocchi, W., Zignani, M., Zollo, F.: Public discourse and social media: a quantitative, cross-platform analysis of the Italian referendum. In: Proceedings of International Conference on Data Science and Advanced Analytics (2017). (to appear)
20. Godbole, N., Srinivasaiah, M., Skiena, S.: Large-scale sentiment analysis for news and blogs. *ICWSM* **7**(21), 219–222 (2007)
21. Hogenboom, A., Bal, D., Frasinca, F., Bal, M., de Jong, F., Kaymak, U.: Exploiting emoticons in sentiment analysis. In: Proceedings of ACM SAC (2013)
22. Stieglitz, S., Dang-Xuan, L.: Emotions and information diffusion in social media—sentiment of microblogs and sharing behavior. *J. Manag. Inf. Sys.* (2013)
23. Mocanu, D., Rossi, L., Zhang, Q., Karsai, M., Quattrociocchi, W.: Collective attention in the age of (mis)information. *Comput. Hum. Behav.* **51**(B) (2015)
24. Facebook. Using the Graph API. Website (2017)
25. Mozetič, I., Grčar, M., Smailović, J.: Multilingual Twitter sentiment classification: the role of human annotators. *PLoS ONE* **11**(5) (2016)
26. Krippendorff, K.: Content Analysis, An Introduction to its Methodology (2012)
27. Vapnik, V.N.: The Nature of Statistical Learning Theory (1995)
28. Kloumann, I.M.C., Danforth, M., Harris, K.D., Bliss, C.A., Dodds, P.S.: Positivity of the English language. *PLoS ONE* **7**(1) (2012)
29. Dodds, P.S., Clark, E.M., Desu, S., Frank, M.R., Reagan, A.J., Williams, J.R., Mitchell, L., Harris, K.D., Kloumann, I.M., Bagrow, J.P., Megerdumian, K., McMahon, M.T., Tivnan, B.F., Danforth, C.M.: Human language reveals a universal positivity bias. *PNAS* **112**(8) (2015)
30. Kralj Novak, P., Smailović, J., Sluban, B., Mozetič, I.: Sentiment of emojis. *PLoS ONE* **10**(12) (2015)