Controversial information spreads faster and further than non-controversial information in Reddit

Jasser Jasser¹ · Ivan Garibay² · Steve Scheinert³ · Alexander V. Mantzaris⁴

Received: 27 June 2020 / Accepted: 26 April 2021 / Published online: 25 May 2021
© The Author(s) 2021

Abstract

Online users discuss and converse about all sorts of topics on social networks. Facebook, Twitter, Reddit are among many other networks where users can have this freedom of information sharing. The abundance of information shared over these networks makes them an attractive area for investigating all aspects of human behavior on information dissemination. Among the many interesting behaviors, controversy within social cascades is of high interest to us. It is known that controversy is bound to happen within online discussions. The online social network platform Reddit has the feature to tag comments as controversial if the users have mixed opinions about that comment. The difference between this study and previous attempts at understanding controversy on social networks is that we do not investigate topics that are known to be controversial. On the contrary, we examine typical cascades with comments that the readers deemed to be controversial concerning the matter discussed. This work asks whether controversially initiated information cascades have distinctive characteristics than those not controversial in Reddit. We used data collected from Reddit consisting of around 17 million posts and their corresponding comments related to cybersecurity issues to answer these emerging questions. From the comparative analyses conducted, controversial content travels faster and further from its origin. Understanding this phenomenon would shed light on how users or organization might use it to their help in controlling and spreading a specific beneficiary message.

✉ Ivan Garibay
igaribay@ucf.edu

Jasser Jasser
Jasser.Jasser@ucf.edu

¹ Department of Computer Science, University of Central Florida, Orlando, FL, USA

² Department of Industrial Engineering and Management Systems, University of Central Florida, Orlando, FL, USA

³ Deloitte Consulting, LLC, Lake Mary, FL, USA

⁴ Department of Statistics and Data Science, University of Central Florida, Orlando, FL, USA

Keywords Controversiality · Information · Diffusion · Reddit · Polarization

Introduction

The emergence of online social networks opened the door for many to amplify content freely. Much of this activity is benign or even positive. However, this openness formed a platform for the polarization of opinions and controversial discussions. There is active research into polarization on online social networks [2, 11, 23]. Controversiality of content is one feature that could draw attention to that content [15]. The spreading and promotion of controversial content support the polarization of social networks, especially in the political sphere [6]. A question naturally arises of whether such behavior is beneficial to the content's promoter and that there is an incentive for the activity. This work investigates this question and whether controversy within a discussion would bring more user attention to authored content and promote the spread of the material further and faster.

Reddit is a popular online platform that allows users to spread knowledge and share opinions through posts consisting of both textual and visual elements. It allows users to respond to others' posts by either commenting or voting for or against other users' posts and comments, known as up- and down-votes [21]. Social media users display a tendency to join homogeneous communities that share the same beliefs and are dedicated to the same topic (echo chambers [12, 22]). Reddit provides spaces, known as *subreddits*, where these discussions occur and which facilitate more focused discussion, and the presence of controversiality is genuine and not artificial. A study by the Pew Research Center (2013) concluded that 6% of online adults are Reddit users [9], and by the time of writing, Alexa¹ ranks it as #6 in the United States and #15 globally². Reddit uses up- and down-votes to identify controversial content; with these votes representing agreement or disagreement, respectively, if the comment received a fair amount of polarized votes, then the comment is considered controversial. The identification of controversial content is determined by the Reddit's formula for controversial comments, the final labeling for which is provided through the Reddit API. According to Reddit's definition, a comment is controversial if 1) the sum of up- and down-votes is greater than or equal to 7, and 2) the ratio of up- and down-votes lies between selected upper ($ub = 0.6$) and lower bounds ($lb = 0.4$) as shown in equation (1). We found the algorithm responsible for identifying controversial comments in the *reddit-archive* repository in GitHub³.

$$C_i = \begin{cases} \text{True}, & lb \leq c \leq ub \\ \text{False}, & \text{else} \end{cases} \quad (1)$$

¹ A web traffic analysis company owned by Amazon

² <https://www.alexa.com/siteinfo/reddit.com>.

³ <https://github.com/reddit-archive/reddit/blob/master/r2/r2/models/builder.py>

Here C_i is the comment index and whether it is classified as controversial or not, while c is either the up- or down-votes ratio. Additionally, *subreddit* moderators may tag a post or comment as controversial, regardless of the number of votes. We assume that the tagging of a comment to be controversial is the result of users being invested in the discussion and that the polarization in the voting comes as users are swayed by some aspects of the discussion.

We rely on the definition and labeling of controversiality that Reddit applies to comments made on the platform. We understand the label to mean that there was attention to what is being discussed and there was a polarized reaction to the labeled comment. The measure of the effect of the resurgence of such labeled comments is studied by analyzing the characteristics of the posts' cascades that contain these comments and those that do not, and by investigating the independent authors' networks that forms these posts' cascades. The results show that conversations that contain controversial content during the early stages of their lifetime possess more activity, and users involved in that conversation have a wider influence than conversations that are not labeled as controversial.

Related work

The concept of polarization and controversiality has been explored in multiple fields within the literature, from its identification in social networks and blogs [1, 3, 6] to its detection in online news and web pages [5, 7, 14]. There are two main themes of the research, the *identification* and *quantification* of controversy within the discussion. Other literature explores the effect of controversial information such as the work of [11], which studies the effect of collective attention on the evolution of controversial debates. The work of [18] explores the effect of the controversy on the emotions and language within online news.

The identification of controversy in social discussions is an area of on-going research with many different approaches to the identification of controversiality. The work of [5] identifies controversy within a topic by looking at the sentiment of the discussion and the range of sentiment polarity across the users involved in the controversy, as measured through natural language processing and sentiment analysis. Similarly in the work of [19], the authors proposed identifying controversial events by also considering the measurement of polarized opinions in the wake of a widely-viewed event. The work of [20] investigated opinion distribution in social media to identify controversial discussions and the size of their opinion groups.

Case studies that have attempted to quantify controversy and controversiality have focused on techniques and methods to measure the amount of controversy within a network or discussion. The work of [11] quantifies controversy in online echo chambers by considering a topic that is controversial and building the conversation graph of that topic which depicts the opinion alignment between the users. The authors then partition the graph to identify the sides of the controversy where the amount of controversy is measured from the characteristics of that graph. The work of [13], compares two types of social networks, those who are formed from polarized context and those who are not by analyzing the boundary between a pair

of communities. The authors then distinguish between the polarized from the non-polarized communities by the concentration of high-degree nodes found in the boundary. The authors found that polarized networks tend to have low concentration “popular (high-degree)” nodes along the boundary between communities.

Data and Investigation

The data used for this investigation is collected from Reddit and provided as part of the “Computational Simulation of Online Social Behavior (SocialSim)” DARPA program⁴. This vast dataset consists of more than 36 million comments on posts related to cybersecurity issues. For this study, we considered the posts in our dataset that have at least 100 comment. The number of posts collected is 47,940 where the total comments in those posts exceeded 17 million comments. the posts are then analyzed to be classified either as controversially initiated or non-controversial based on whether there exist any comments that are labeled controversial according to Reddit (as explained above) within the early stages of the posts. The classification yielded 23,101 posts that are labeled as controversially initiated, and 24,839 are labeled non-controversial. The distribution of controversial and non-controversial cascades from the same *subreddits* is almost equal in quantity. As an example, there are 8,528 controversial and 12,349 non-controversial cascades coming from the ‘*pcmasterrace*’ *subreddit*, and 5,450 controversial and 4,149 non-controversial cascades coming from the ‘*android*’ *subreddit*.

In social networks such as Reddit, users publish posts in regard of a specific topic where they share information and opinions. The posts are time stamped and so are the comments from the users that replied to that post. This forms a temporal propagation graph, also known as information cascade [4, 17]. The classification of cascades is based on the enclosure of controversially labeled comments within the posts that form the cascade. Cascades have a growth stage which is the period between the time of the first comment (t_0) and the peak time (t_p) of the post’s cascade as shown in Fig. 1a. We define a cascade that follows a post as controversial if the cascade contains a comment, directed to the post’s author, that is labeled controversial and occurred within the growth stage. Assuming that the discussion reached a level of polarity in the opinions within the growth stage of the cascade that drove the discussion. We do not consider controversial comments that happened after the growth stage since at that stage users have lost interest in the topic and the polarization did not reinvigorate the discussion. Fig. 1 shows posts’ cascades where Fig. 1a–c are classified as controversial while Fig. 1d–f are classified as non-controversial. Notice Fig. 1e,f where in Fig. 1e there are two controversial comments after its growth stage and for that it was classified as non-controversial, and in Fig. 1f there is one controversial comment that is within the growth stage, however, it was not directed to the post’s author and so it was not classified as controversial.

⁴ <https://www.darpa.mil/program/computational-simulation-of-online-social-behavior>.

Fig. 1 Controversially initiated and non-controversially initiated cascades, **a–c** are controversially initiated posts’ cascades while **d–f** are non-controversial posts’ cascades where the red dots represent a comment labeled as controversial by Reddit that is directed to the post’s author while a green dot is a comment labeled controversial by Reddit that is directed to another comment

In our assumption, when the Reddit algorithm labels a comment as controversial, it is a sign of collective attention [24] where users taking part in that discussion are still active and paying attention. The attention in large groups to the topics indicates popularity, where the popular the topic, the faster and further the information sharing and dissemination [24]. The comments directed to the post author are more visible to the reader, thus poised to grab more attention, and for that, we prioritized them in our analyses. As an example, the largest cascade in our dataset (controversially initialized) is about a giveaway for a personal computer. Comments directed to the post’s author such as “Give it to meeeh Im poor,” and “Please for the love of God please pick me I never win anything” got attention with polarized responses and were labeled controversial by Reddit’s algorithm. The attention within this cascade drove it to end up having around 60,000 comments. While the largest non-controversial cascade was also concerning a personal computer giveaway, but it did not hold any controversially labeled comments directed to the post’s author and did not achieve the high amount of comments (around 23,000 comments) mentioned in the previous cascade. However, there are controversially labeled comments that are not directed to the post’s author and occurred after the growth stage of the cascade where the attention and discussion have decayed.

The data come with the sentiment polarity and subjectivity of the comments calculated using the “pattern.en” model⁵. The model ranks the polarity of the comment sentiment based on the contained adjectives on a scale from -1 to 1 where -1 means

⁵ <https://www.clips.uantwerpen.be/clips.bak/pages/pattern-en#sentiment>.

Fig. 2 Authors' averages of sentiment polarity and subjectivity, **a** shows the authors' average for sentiment polarity while **b** shows the authors' average for sentiment subjectivity

highly negative sentiment while 1 is a highly positive sentiment. The subjectivity is ranked by the model on a scale from 0 to 1 where 0 is very objective and 1 is very subjective. The sentiment polarity and subjectivity of the authors were calculated by taking the average sentiment of all comments written by that author. If an author wrote a comment that got labeled as controversial by Reddit, that author is labeled as controversial. The dataset contains 1,172,886 unique authors. Controversial authors consist of around 6% of that total or 69,607 authors to be precise. Figure 2 shows that authors that are labeled controversial tend to have balanced polarity that is normally distributed with a mean slightly above 0, meaning they are more likely to be positive in their comments. Unlike the controversial authors, non-controversial authors mostly have no sort of polarity in their text, putting a spike in their distribution at 0. However, like controversial authors, there is a smaller second peak that is slightly positive suggesting a bi-modal distribution. In terms of subjectivity, controversial authors tend to have a balance between being objective or subjective about the topic, hence their sentiment subjectivity average is closer to 0.5 while non-controversial authors have a normal distribution with a larger variance, but their peak is almost half the size of controversial authors' distribution, and there is a sharp peak at 0 indicating more objectivity from that distribution.

A social network is a form of complex network [10], a social structure that consists of social actors and the relationships between them. The relationship and connectedness between the actors define and distinguish those that are influencing the network from those who do not. The work of [16], argues that most influencers are found within the k -core of the network. A k -core is the "largest subgraph where vertices have at least k interconnections" [8]. We considered the independent authors that commented on each other comments for every post to measure the connectedness among them and discovering the subset of nodes with the highest *coreness*. This subset of authors defines the nodes responsible for influencing the dissemination of information. The k -core of that graph network is analyzed where k is $\max(\text{coreness})$. We are interested in finding the number of nodes that are labeled controversial and their ratio within these k -cores.

Fig. 3 Descriptive analysis results of posts' cascade, **a** shows the difference in size between controversially initiated post's cascade and non-controversial one. **b** shows the difference in peak size, and **c** shows the difference in the number of unique authors

We investigated whether the earlier the resurgence of the controversial comments within the cascade would lead to drastic changes in the characteristics by defining epochs within the growth stage of the cascade. The period is then divided into three epochs, epoch 1 for the first quartile, epoch 2 for the second quartile, and epoch 3 for the second half of the growth stage. If one of these epochs contains the majority of the controversial comments, we consider that epoch as the period with the highest concentration of controversy, and we associate the post's cascade with that period. An epoch 1 cascade means that the cascade's highest concentration of controversy happened during the first quartile of the growth stage.

Results

The results of the analysis show that posts' cascades that contained controversially labeled messages during its growth stage produce a larger total number of comments and a larger burst, indicating more involvement from users. Another feature is that

Fig. 4 Network analysis results of posts' cascade, **a** shows the difference in network size (number of edges) between controversially initiated post's cascade and non-controversial one, **b** shows the difference in the number of nodes within the networks' k -core, and **c** shows the difference in the number of controversial nodes within the networks' k -core

more unique users are involved in the cascades which is controversially labeled. This clearly shows that its beneficial for a user who seeks to spread a particular message faster and further, to infuse controversy within the context of the message. The plots in Fig. 3 show this effect presented on three different measurements. Each plot is a CCDF (complementary cumulative distribution function) where the x -axis is the quantity in question shown in the title and the y -axis is the probability for that measure to exceed that x -axis value. In Fig. 3a, shows that the posts that attract the largest contributions are dominated by initial signals of controversy. As the cascades grew, a difference emerged; non-controversial cascades did not exceed 25,000 comments whereas the controversial cascades did display posts having 60,000 or more comments. Figure 3b shows the peak size probability where controversial cascades tend to generally have larger peaks. Figure 3c shows the probability for cascades to exceed a certain number of unique authors contributing comments to a post. Along with this measure it can be seen how the controversially initiated cascades produce much larger values for the upper tail of the distribution.

Fig. 5 Network evolution analysis results of posts' cascade, **a** shows the difference in network size (number of edges) hourly growth between controversially initiated post's cascade and non-controversial one. **b** Shows the difference in the hourly growth of the number of nodes within the networks' k -core, and **c** shows the difference in the hourly growth of the number of controversial nodes within the networks' k -core

Figure 4 shows the results from conducting a network analysis on the post structure in the Reddit dataset. The results here show a more distinctive differentiation between the two types of cascades. Figure 4a shows the CCDF for the network size of the controversially and non-controversially initiated cascades where the network size is computed by using the total number of edges produced within the entity of the post. Figure 4b looks at the number of nodes within the defined k -core in the cascades' network. It is clearly visible that controversial cascades end up containing more nodes within their k -core and so are the controversial nodes within these k -cores as shown in Fig. 4c. This shows that network properties of the different cascades are more affected by content which attracts controversial initiations or that they shape the subsequent discussions.

Figure 5 shows an evolutionary temporal analysis of the network changes within the first 24 h. Figure 5a shows the evolution the number of links in the network. The activity for the controversial cascades is greater and for the time duration observed has this increase from each in the post creation. Figure 5b shows the k -core evolution

Fig. 6 The posts' cascades size and peaks size based on the epoch that contains the highest concentration of the controversial comments. **a** Shows the comparison of posts' cascades size, while **b** shows the comparison of the peak size

where the number of nodes within the controversial k -core increased drastically over the other ones. Figure 5c compares the evolutionary number of controversial nodes found within the k -core. This reinforces that the cascades which contain controversial activity during their early stages continue to swiftly progress for longer periods of time than the non-controversial cascades.

Figure 6 shows the results of comparing the first three epochs by the cascade size and the peak size of the cascade. Figure 6a shows that epoch 1 controversial cascades produced the largest number of comments and peak size, followed by epoch 2 then epoch 3 cascades. This indicates that the earlier the resurgence of the controversiality within the cascade the further the information will disseminate within, while later resurgence of controversial comments might not drastically affect the information flow.

Conclusion

The work sought to explore whether controversial content will have a greater chance of increased activity amongst other users. The analysis uses descriptive statistics and network analysis for the amount of activity a cascade produced. The k -core analysis conducted to study the structure of the networks' core. The temporal dimension is also explored in an analysis of the evolution of contributors. The results show that content which was controversial is associated with higher degrees of activity. This can shed light on various marketing strategies for competing for attention online. Even if not completely understood as to why this occurs, it can be utilized by accounts on social media seeking to create more influence. It can also help explain the spread of disinformation.

Acknowledgements This work was partially supported by grant FA8650-18-C-7823 from the Defense Advanced Research Projects Agency (DARPA). The views and opinions contained in this article are the

authors and should not be construed as official or as reflecting the views of the University of Central Florida, DARPA, or the U.S. Department of Defense.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Adamic, L. A., & Glance, N. (2005). The political blogosphere and the 2004 u.s. election: Divided they blog. In Proceedings of the 3rd International Workshop on Link Discovery, LinkKDD '05, pp. 36–43. ACM, New York, NY, USA. <https://doi.org/10.1145/1134271.1134277>. <http://doi.acm.org/10.1145/1134271.1134277>.
2. Akoglu, L. (2014). Quantifying political polarity based on bipartite opinion networks. In Eighth International AAAI Conference on Weblogs and Social Media.
3. An, J., Quercia, D., & Crowcroft, J. (2014). Partisan sharing: Facebook evidence and societal consequences. In Proceedings of the Second ACM Conference on Online Social Networks, COSN '14, pp. 13–24. ACM, New York, NY, USA. <https://doi.org/10.1145/2660460.2660469>. <http://doi.acm.org/10.1145/2660460.2660469>.
4. Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5), 992–1026.
5. Choi, Y., Jung, Y., & Myaeng, S. H. (2010). Identifying controversial issues and their sub-topics in news articles. In Pacific-Asia Workshop on Intelligence and Security Informatics, pp. 140–153. Springer.
6. Conover, M. D., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F., & Flammini, A. (2011). Political polarization on twitter. In Fifth international AAAI conference on weblogs and social media.
7. Dori-Hacohen, S., & Allan, J. (2013). Detecting controversy on the web. In Proceedings of the 22nd ACM international conference on Conference on information & knowledge management, CIKM '13, pp. 1845–1848. ACM, New York, NY, USA. <https://doi.org/10.1145/2505515.2507877>. <http://doi.acm.org/10.1145/2505515.2507877>.
8. Dorogovtsev, S. N., Goltsev, A. V., & Mendes, J. F. F. (2006). K-core organization of complex networks. *Physical Review Letters*, 96(4), 040601.
9. Duggan, M., & Smith, A. (2013). 6% of online adults are reddit users. *Pew Internet & American Life Project*, 3, 1–10.
10. Easley, D., Kleinberg, J., et al. (2010). *Networks, crowds, and markets* (Vol. 8). Cambridge University Press Cambridge.
11. Garimella, K., Morales, G. D. F., Gionis, A., & Mathioudakis, M. (2018). Quantifying controversy on social media. *ACM Transactions on Social Computing*, 1(1), 3.
12. Garrett, R. K. (2009). Echo chambers online? Politically motivated selective exposure among internet news users. *Journal of Computer-Mediated Communication*, 14(2), 265–285.
13. Guerra, P. C., Meira Jr, W., Cardie, C., & Kleinberg, R. (2013). A measure of polarization on social media networks based on community boundaries. In Seventh International AAAI Conference on Weblogs and Social Media.

14. Kaplun, K., Leberknight, C., & Feldman, A. (2018). Controversy and sentiment: An exploratory study. In Proceedings of the 10th Hellenic Conference on Artificial Intelligence, SETN '18, pp. 37:1–37:7. ACM, New York, NY, USA. <https://doi.org/10.1145/3200947.3201016>. <http://doi.acm.org/10.1145/3200947.3201016>.
15. Kim, J., Wyatt, R. O., & Katz, E. (1999). News, talk, opinion, participation: The part played by conversation in deliberative democracy. *Political Communication*, 16(4), 361–385.
16. Kitsak, M., Gallos, L. K., Havlin, S., Liljeros, F., Muchnik, L., Stanley, H. E., et al. (2010). Identification of influential spreaders in complex networks. *Nature Physics*, 6(11), 888.
17. Leskovec, J., Krause, A., Guestrin, C., Faloutsos, C., Faloutsos, C., VanBriesen, J., & Glance, N. (2007). Cost-effective outbreak detection in networks. In Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 420–429. ACM.
18. Mejova, Y., Zhang, A. X., Diakopoulos, N., & Castillo, C. (2014). Controversy and sentiment in online news. arXiv preprint [arXiv:1409.8152](https://arxiv.org/abs/1409.8152).
19. Popescu, A. M., & Pennacchiotti, M. (2010). Detecting controversial events from twitter. In Proceedings of the 19th ACM International Conference on Information and Knowledge Management, CIKM '10, pp. 1873–1876. ACM, New York, NY, USA. <https://doi.org/10.1145/1871437.1871751>. <http://doi.acm.org/10.1145/1871437.1871751>.
20. Qiu, J., Lin, Z., & Shuai, Q. (2019). Investigating the opinions distribution in the controversy on social media. *Information Sciences*, 489, 274–288.
21. Stoddard, G. (2015). Popularity dynamics and intrinsic quality in reddit and hacker news. In Ninth International AAAI Conference on Web and Social Media.
22. Sunstein, C. R. (2001). *Echo chambers: Bush V. Gore, impeachment, and beyond*. Princeton, NJ: Princeton University Press.
23. Taylor, C., Mantzaris, A., & Garibay, I. (2018). Exploring how homophily and accessibility can facilitate polarization in social networks. *Information*, 9(12), 325.
24. Wu, F., & Huberman, B. A. (2007). Novelty and collective attention. *Proceedings of the National Academy of Sciences*, 104(45), 17599–17601. <https://doi.org/10.1073/pnas.0704916104>. <https://www.pnas.org/content/104/45/17599>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.