

# Echo-chambers and Idea Labs: Communication Styles on Twitter

Aleksandra Sorokovikova, Michael Becker, Ivan P. Yamshchikov

Constructor University, Bremen

Technical University of Applied Sciences Würzburg-Schweinfurt (THWS), Würzburg  
CAIRO, Technical University of Applied Sciences Würzburg-Schweinfurt (THWS), Würzburg  
ivan.yamshchikov@thws.de

## Abstract

This paper investigates the communication styles and structures of Twitter (X) communities within the vaccination context. While mainstream research primarily focuses on the echo-chamber phenomenon, wherein certain ideas are reinforced and participants are isolated from opposing opinions, this study reveals the presence of diverse communication styles across various communities. In addition to the communities exhibiting echo-chamber behavior, this research uncovers communities with distinct communication patterns. By shedding light on the nuanced nature of communication within social networks, this study emphasizes the significance of understanding the diversity of perspectives within online communities.

**Keywords:** echo chambers, computational social science, idea labs

## 1. Introduction

Online social environments are often characterised by the phenomenon of the so-called echo chambers where participants isolate themselves from opposing opinions, reinforcing their own beliefs through limited communication within their community. These chambers can lead to ad hominem attacks, targeting individuals rather than engaging with their arguments, and straw man arguments, which distort opposing viewpoints for easier dismissal. These dynamics contribute to polarization and the adoption of more extreme positions (Petit et al., 2020). Polarization and echo chambers are commonly observed in social networks, facilitated by recommendation algorithms (Shore et al., 2018; Rossi et al., 2018).

To explore these communities, various algorithms such as the Infomap (Rosvall and Bergstrom, 2008), Louvain algorithm (Campigotto et al., 2014), Stochastic Blockmodels (e.g. (Peixoto, 2020)) or force-directed layout (Gaisbauer et al., 2023) are employed to identify clusters within the user interaction graph. In such setting, users are represented as vertices, and interactions such as social network connections, retweets, and replies on platforms like Twitter (X) are captured as graph edges. These graphs are then partitioned into dense clusters, interpreted as communities sharing similar opinions or engaging in similar activities. Understanding the structure of these communities involves analyzing not only user interactions but also the content they generate. Content analysis allows us to capture the characteristics of the produced content itself (Garimella et al., 2016). Previous studies have investigated echo chambers and polarization in social media, particularly concerning topics like

COVID-19, proposing models such as RetweetBERT and DICE for sentiment analysis and detecting ad hominem attacks (Jiang et al., 2021; Naseem et al., 2020; Delobelle et al., 2019).

In this paper, we focus on identifying Twitter (X) communities related to vaccination. We employ community detection algorithms to identify clusters based on user interactions and the content of their tweets. Additionally, we train classifiers for content analysis, such as sentiment, subjectivity, ad hominem, and straw man arguments. Using these classifiers, we evaluate communication style that characterises each community. This approach enables us to uncover that the patterns of user interaction within communities are clearly different. Moreover, to some extent one could identify community of the user based on their communication style. Thus, we suggest that it is sub-optimal to lump every community under a broad umbrella term "echo-chamber". Instead, we suggest there is a need for a more detailed taxonomy based on the detected systematic differences in the communication styles.

## 2. Data

For this work a ready-made twitter (X) vaccination dataset<sup>1</sup> was taken, it contains approximately 1.5m tweets on the vaccination topic and 770k unique users. This dataset was collected with TWINT - open-source scraping tool. This dataset is suitable for research for several reasons. Firstly, all tweets relate to the topic of vaccination, and the tweets are taken from a very wide time range (starting in

<sup>1</sup><https://www.kaggle.com/datasets/keplaxo/twitter-vaccination-dataset>

2006, ending in 2019). Secondly, it is possible to build a reply graph, since for tweets that are reply, there is the id of the user to whose tweet this reply was made. A large number of tweets collected in the dataset allows to cover the topic of vaccination from different sides and opinions. Finally, all the discussions are thematically aligned so the differences between texts written by the members of different communities are less prone to be topically aligned. All the discussions are centered around one general topics and the stylistic differences between the texts are more differentiating than the topics these texts discuss. This makes the dataset a good case-study for the core hypothesis of stylistic distinctions between communications styles of various communities.




Figure 1: Conversation graph after applying OpenOrd algorithm

### 3. Experiments

We calculate a variety of metrics using texts of the tweets without any information on the user community in which a given tweet is published. Then we cluster the tweets based on those classifiers and compare mutual information between the obtained text-based clustering and the clusters that are formed in the communication graph. It turns out that the textual information alone partly allows to predict the community in which the user communicates.

#### 3.1. Community clustering

Based on the dataset, a graph was built as follows: each vertex represents a user, and an edge is created if user A made a reply to user B. To get rid of

noise and make the graph more dense, a weight was added to the edges, which meant the number of replays of the user. Only edges with weight 3 or higher remained in the final graph. Subsequently, the OpenOrd algorithm (Martin et al., 2011) was used for spatialization, arranging the vertices in two-dimensional space with  $x$  and  $y$  coordinates, as depicted in Figure 1. Notably, for community detection, the Louvain algorithm was utilized, with clusters determined based on modularity. The computed modularity value for this graph was 0.902, and the subsequent analysis focused on six of the densest communities, informed by clustering coefficient and modularity considerations.

#### 3.2. Text Classification Metrics

We calculate a variety of metrics that assess communication style and are based solely on the textual content of the tweets.

##### 3.2.1. Polarity Scores

The polarity scores were calculated using the Optimized BERT Pretraining Approach (RoBERTa), which was trained on around 58 million tweets. We used the pretrained RoBERTa-based classifier developed in (Barbieri et al., 2020) to calculate the negativity, neutrality and positivity of a give textual input.

##### 3.2.2. Subjectivity

The subjectivity score was determined using the TextBlob (Loria et al., 2018) library. The subjectivity value in TextBlob indicates the degree of subjective or objective content of a given text. Subjectivity refers to how opinionated or subjective the text is, while objectivity refers to a more factual or objective writing style. The subjectivity value is a floating point value between 0.0 and 1.0, where 0.0 indicates a very objective or factual text and 1.0 indicates a very subjective or opinionated text.

##### 3.2.3. Logical fallacy

The two values "label" and "score" were determined using a pre-trained model for logical fallacy detection (Jin et al., 2022). The score label can take values between 0 and 12 and encodes various logical fallacies, namely: Ad Hominem, Ad Populum, False Causality, Circular Claim, Appeal to Emotion, Fallacy of Relevance, Deductive Fallacy, Intentional Fallacy, Fallacy of Extension, False Dilemma, Fallacy of Credibility, Equivocation. The model assigns a probability between 0 and 1 for every given label.

text/user based cluster	1	2	3	4	5	6
1	445**	942*	172	128	70	58
2	720**	2763*	413	323	239	156
3	262	974*	355**	71	111	71
4	237	827*	94	264**	86	64
5	202	736*	162	82	308**	41
6	205	718*	145	92	71	492**

Table 1: Confusion matrix for text-based clustering used for user-based cluster prediction. The clusters are unbalanced. Cluster number 2 has the biggest support, thus \* marks the highest value in a given row, while \*\* marks the second highest

text/user based cluster	1	3	4	5	6
1	<b>445</b>	172	128	70	58
3	262	<b>355</b>	71	111	71
4	237	94	<b>264</b>	86	64
5	202	162	82	<b>308</b>	41
6	205	145	92	71	<b>492</b>

Table 2: Confusion matrix for text-based clustering used for interaction-based cluster prediction. Cluster number 2 is removed. The highest values in a given row are marked bold.

#### 4. Predicting Community Membership with Communication Style

Now every tweet could be described by a set of various classifier scores. At the same time we know the community to which the author of the tweet belongs since we have the structure of the reply clusters that we obtained in Section 3.1. To test whether every echo-chamber is characterised with similar communication style we can build clusters of tweets based solely on the classifier scores. Normally, a choice for the number of clusters in a clustering could be difficult. However, since the reply graph clustering procedure has already detected six clusters we can choose six clusters for our text-based clustering as well. Now every tweet belongs to one text-based cluster as well as one reply graph based cluster. Comparing those labels allows us to see to which extent solely textual content of the tweet informs us on the community in which given communication occurs.

Figure 2 demonstrates average scores for six clusters detected in Section 3.1. One can immediately see that all six are characterized by rather different communication styles. One can see that cluster number two is represented by the point in the "center". It is the biggest cluster that includes a lot of weakly connected users that do not form a dense clique. This cluster has the biggest support in terms of absolute numbers but represents users who are

not active members of any one of the five dense communities but are rather occasional posters. Thus, it stands to reason that the average scores of the classifiers for the tweets in this cluster end up in the center of the cloud of points. Table 1 shows the confusion matrix between knn-clusters based on texts of the tweets and the clusters obtained from the graph of interactions between users. The accuracy of the user-based clustering when predicted by text is 0.35 which is quite interesting in itself.

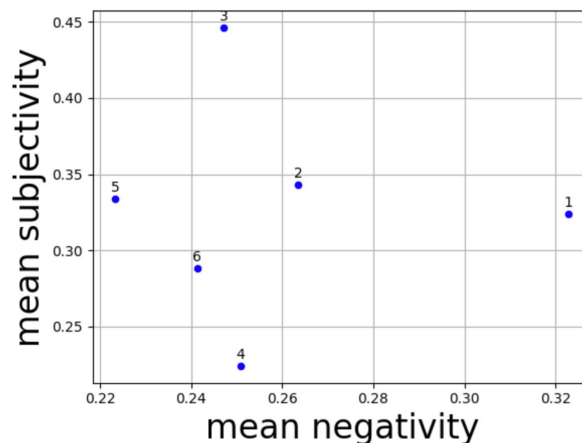


Figure 2: Average scores for mean subjectivity and mean negativity in all 6 user communities. One can clearly see that one of the communities is characterized as highly subjective while another is highly negative while scoring lower on subjectivity.

If we drop this cluster number two that represents weakly connected users that do not form a tight community the accuracy reaches 0.44, see Table 2. This highlights that tight communities have a distinct expressive communication style indeed.

#### 5. Discussion

This study offers a case study and an initial exploration of the varied communication styles within Twitter (X) communities discussing vaccination, challenging the simplistic notion of echo chambers. Through the application of community detection al-

gorithms and text classification metrics, we demonstrate significant diversity in subjectivity, negativity, and logical fallacies across different groups. This suggests a broader spectrum of communicative behaviors in online discussions than typically discussed.

Our findings highlight the importance of nuanced understanding of online discourse, pointing towards the necessity for a more detailed taxonomy of communication styles. This work suggests a pathway for future research to explore the complexities of digital community communication, advocating for a deeper examination beyond conventional categorizations to better understand how these communities interact and evolve.

## 6. Conclusion

In our investigation of communication patterns within Twitter's (X's) echo chambers, we aimed to identify variations in discourse, specifically seeking environments akin to 'Idea Labs' where open, critical discussion prevails over personal attacks. Utilizing computational methods, we analyzed a substantial dataset to discern these communication styles.

Our results did not confirm the presence of 'Idea Labs' in the studied dataset. However, the study revealed significant variations in communication styles across different echo chambers. Despite discussing identical topics, the textual characteristics within each community were distinct enough to allow for a predictive model to accurately categorize tweets based on their origin.

This finding is critical, demonstrating the extent to which echo chambers can influence discourse style. It also highlights the potential for computational approaches to identify and categorize such patterns in online communication.

## Limitations

This paper focuses on one particular case-study. The topic of the discussions is specific and all the results are limited to twitter (X) discussions only. Thus one can not be sure that the results are general and could be applicable to other social media communities or to the discussions around other topics. However, the provided case-study is a good starting point to initiate a deeper discussion of a more nuances approach to community formation in social media that regards the phenomenon in a more holistic manner and takes into account both the structure of the social graph as well as the content of the communications.

## Acknowledgements

The authors are deeply thankful to Dr. Felix Gaisbauer for his support, fruitful discussions and valuable advice.

## A. Bibliographical References

- Alfred V. Aho and Jeffrey D. Ullman. 1972. *The Theory of Parsing, Translation and Compiling*, volume 1. Prentice-Hall, Englewood Cliffs, NJ.
- American Psychological Association. 1983. *Publications Manual*. American Psychological Association, Washington, DC.
- Rie Kubota Ando and Tong Zhang. 2005. [A framework for learning predictive structures from multiple tasks and unlabeled data](#). *Journal of Machine Learning Research*, 6:1817–1853.
- Galen Andrew and Jianfeng Gao. 2007. [Scalable training of  \$L\_1\$ -regularized log-linear models](#). In *Proceedings of the 24th International Conference on Machine Learning*, pages 33–40.
- Francesco Barbieri, Jose Camacho-Collados, Leonardo Neves, and Luis Espinosa-Anke. 2020. Tweeteval: Unified benchmark and comparative evaluation for tweet classification. *arXiv preprint arXiv:2010.12421*.
- Romain Campigotto, Patricia Conde-Céspedes, and Jean-Loup Guillaume. 2014. A generalized and adaptive method for community detection.
- Ashok K. Chandra, Dexter C. Kozen, and Larry J. Stockmeyer. 1981. [Alternation](#). *Journal of the Association for Computing Machinery*, 28(1):114–133.
- James W. Cooley and John W. Tukey. 1965. [An algorithm for the machine calculation of complex Fourier series](#). *Mathematics of Computation*, 19(90):297–301.
- Pieter Delobelle, Murilo Cunha, Eric Cano, Jeroen Peperkamp, and Bettina Berendt. 2019. [Computational ad hominem detection](#).
- Felix Gaisbauer, Armin Pournaki, Sven Banisch, and Eckehard Olbrich. 2023. Grounding force-directed network layouts with latent space models. *Journal of Computational Social Science*, pages 1–33.
- Kiran Garimella, Gianmarco Morales, Aristides Giornis, and Michael Mathioudakis. 2016. [Quantifying controversy in social media](#).



- Dan Gusfield. 1997. *Algorithms on Strings, Trees and Sequences*. Cambridge University Press, Cambridge, UK.
- Mathieu Jacomy, Tommaso Venturini, Sebastien Heymann, and Mathieu Bastian. 2014. [Forceatlas2, a continuous graph layout algorithm for handy network visualization designed for the gephi software](#).
- Julie Jiang, Xiang Ren, and Emilio Ferrara. 2021. [Social media polarization and echo chambers: A case study of covid-19](#).
- Zhijing Jin, Abhinav Lalwani, Tejas Vaidhya, Xiaoyu Shen, Yiwen Ding, Zhiheng Lyu, Mrinmaya Sachan, Rada Mihalcea, and Bernhard Schölkopf. 2022. [Logical fallacy detection](#).
- Steven Loria et al. 2018. textblob documentation. *Release 0.15*, 2(8):269.
- Shawn Martin, W. Brown, Richard Klavans, and Kevin Boyack. 2011. [Openord: An open-source toolbox for large graph layout](#).
- Usman Naseem, Imran Razzak, Katarzyna Musial, and Muhammad Imran. 2020. [Transformer based deep intelligent contextual embedding for twitter sentiment analysis](#).
- Tiago P Peixoto. 2020. Latent poisson models for networks with heterogeneous density. *Physical Review E*, 102(1):012309.
- John Petit, Cong Li, and Khudejah Ali. 2020. [Fewer people, more flames: How pre-existing beliefs and volume of negative comments impact online news readers' verbal aggression](#).
- Mohammad Sadeq Rasooli and Joel R. Tetreault. 2015. [Yara parser: A fast and accurate dependency parser](#). *Computing Research Repository*, arXiv:1503.06733. Version 2.
- Wilbert Samuel Rossi, Jan Polderman, and Paolo Frasca. 2018. [The closed loop between opinion formation and personalised recommendations](#).
- Martin Rosvall and Carl T Bergstrom. 2008. Maps of random walks on complex networks reveal community structure. *Proceedings of the national academy of sciences*, 105(4):1118–1123.
- Jesse Shore, Jiye Baek, and Chrysanthos Dellarocas. 2018. [Network structure and patterns of information diversity on twitter](#).