Inside Echo Chambers: An Exploration of Polarization and Topic Diversity on Social Media

Faisal Alatawi, Paras Sheth, and Huan Liu

Arizona State University
{faalataw,psheth5,huanliu}@asu.edu

Abstract. This study explores the intricate dynamics of echo chambers within social media platforms, specifically investigating the impact of polarization levels on the diversity of topics discussed by users. Rather than solely focusing on proving the existence of echo chambers, we delve into their characteristics, user interactions, and impacts on discourse. Leveraging a recently proposed measure, the "Echo Chamber Score" (ECS), we quantify the level of polarization within these chambers, concentrating on communities with both the highest and lowest polarization levels. Our findings reveal a significant correlation between the degree of polarization and the distribution of topics, with highly polarized chambers demonstrating a narrower topical focus. We believe our research will not only pave the way for further studies on echo chambers at a community level, but also aid in devising more effective strategies to promote diversity in online discussions.

Keywords: Echo Chamber · Polarization · Topic Modeling

1 Introduction

In today's digital age, social media platforms have become prominent spaces for information sharing, discussion, and community formation. However, these platforms also have the potential to create echo chambers, where individuals are primarily exposed to ideas and opinions that align with their own, leading to the reinforcement of existing beliefs and limited exposure to diverse perspectives [5]. This phenomenon of polarization and topic homogeneity within echo chambers has raised concerns about the potential negative impacts on public discourse, democratic processes, and societal cohesion [1, 9]. Understanding the dynamics of echo chambers and their influence on topic diversity is crucial for addressing these concerns and developing strategies to foster more inclusive and balanced online environments.

Previous research on echo chambers primarily revolves around detecting their existence and understanding the underlying factors contributing to their formation [5]. A significant focus of these studies lies in examining the ideological interactions among users, particularly in determining whether interacting users share similar or divergent ideological beliefs [6, 3, 5]. However, to the best of our knowledge, there is a notable gap in research concerning the relationship between the diversity of topics that users interact with and the presence of polarization. While some works have attempted to identify communities where users share common topics of interest and exhibit similar sentiments towards these topics [12, 10, 7], our study differentiates itself by investigating the impact of community polarization on the diversity of topics. In contrast, these previous works primarily aim to detect these communities rather than explore the influence of polarization on topic diversity.

The aim of this paper is to delve into the concept of echo chambers and investigate the relationship between polarization and topic diversity on social media. We seek to explore whether the degree of polarization, as measured by the Echo Chamber Score (ECS) [2], impacts the diversity of topics and discussions among users. To accomplish this, we employ a comprehensive methodology that analyzes both dataset-level and community-level topic distributions.

To examine the dataset-level topic distribution, we use four datasets collected by Alatawi at el. [2], including two polarized datasets (Gun and Abortion) and two unpolarized datasets (Super Bowl and SXSW). By employing BERTopic [8] for topic detection and labeling, we identify the prevalent topics within each dataset and observe the variations in topic diversity among the polarized and unpolarized datasets. Our results suggest that individuals who participate in polarized discussions may have a tendency to prioritize political subjects. This observation underscores the possibility of echo chambers within polarized datasets.

Moving beyond dataset-level analysis, we delve into community-level topic distribution using the ECS measure. Focusing on the Abortion (polarized) and SXSW (unpolarized) datasets, we examine the topic compositions within the most polarized and least polarized communities. Our results indicate that more polarized communities tend to engage in more political discussions and display greater topic homogeneity compared to less polarized communities. We hope that this research can informs strategies for promoting a more diverse and balanced discourse on social media platforms. It is important to acknowledge the harmful effects of echo chambers and work towards creating a more inclusive and open online environment, by understanding the complex interplay between polarization and topic diversity, we can better address the challenges of echo chambers and promote a more informed and engaged public discourse.

2 Methodology

In this study, we examine the relationship between echo chambers and topic distribution by analyzing the impact of polarization on the tweets shared by users. The study utilized four datasets collected by Alatawi et al. [2], which consisted of two polarized datasets (Gun and Abortion) and two unpolarized datasets (Super Bowl and SXSW).

To conduct the analysis, the timeline tweets of each user within the selected datasets were used. To ensure data quality, the tweets were preprocessed by retaining only English tweets that contained a minimum of 5 words, excluding short and less informative tweets. This approach to preprocessing mirrored the filtering methodology previously employed by Alatawi et al. [2]. This preprocessing step aimed to focus on meaningful and substantial content for further analysis.

For the analysis of topic distribution, we employed BERTopic [8], a topic modeling technique. BERTopic was used to detect topics in the tweets across the user base, providing insights into the prevalent topics within each dataset. Additionally, GPT-4 [11], a language model, was utilized to provide more informative names for the topic words obtained from BERTopic. This enhanced the interpretability of the identified topics and facilitated a clearer understanding of the discussions within the datasets.

In addition to dataset-level analysis, this paper also focused on studying the distribution of topics at the community level. This involved utilizing the Echo Chamber Score (ECS) [2] as a measure of polarization. Two datasets, the polarized Abortion dataset and the unpolarized SXSW dataset, were selected for this analysis. The Louvain algorithm [4], a popular community detection method, was employed to identify communities within these datasets. Based on the ECS scores, the most polarized and least polarized communities within each dataset were identified for further analysis.

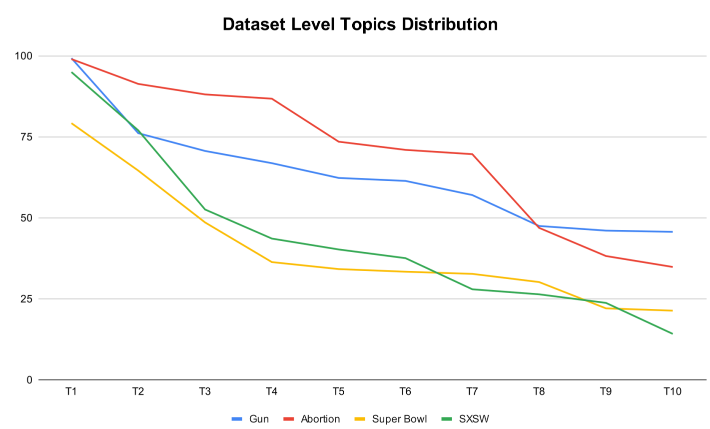

Fig. 1: The distribution of topics across our four datasets. It shows the percentage of users who have tweeted about the top 10 topics in each dataset (see Table 2). Notably, users in the polarized datasets (Gun and Abortion) exhibit higher percentages across all topics, suggesting a greater homogeneity of interests. This finding aligns with the concept of an echo chamber.

3 Experiments

This section focuses on examining the relationship between echo chambers and topic distribution. Specifically, we aim to explore how the level of polarization, as measured by the Echo Chamber Score (ECS), is reflected in the tweets shared by users in their timelines. We hypothesize that users engaged in polarized discussions exhibit less diversity in their shared interests, as manifested by the topics they discuss. Therefore, our research question for this section is as follows: **Does the degree of polarization, as represented by the ECS score, impact the diversity of topics and discussions among users?**

3.1 Impact of Polarization on Dataset-Level Topic Distribution

In this experiment, our goal was to examine the impact of polarization on the conversation surrounding the four datasets collected by Alatawi et al. [2]. We utilized the timeline tweets of each user for analysis. To preprocess the tweets, we retained only English tweets containing a minimum of 5 words to exclude any short, less informative tweets, following Alatawi et al.'s filtering method. Table 1 presents the average number of tweets per user after applying these filters.

Dataset	# Tweets	# Users	Avg Tw/User	# Tweets Assigned topics (%)
Gun	814K	6,894	118	418K (51.3%)
Abortion	768K	5,074	151	386K (50.2%)
Super Bowl	230K	5,015	45	111K (48.3%)
SXSW	252K	2,413	104	129K (51%)

Table 1: The four datasets' information.

Subsequently, we employed BERTopic [8] to detect topics for all the tweets across the user base. Table 1 also displays the number of tweets for which BERTopic successfully assigned a meaningful topic. On average, BERTopic identified topics for 50% of the tweets. This highlights the considerable challenge involved in topic detection for tweets and emphasizes the need for improvements in topic modeling techniques to enhance our understanding of the dataset.

We utilized GPT-4 [11] to provide more informative names for the topic words obtained from BERTopic. The resulting topic names are presented in Table 2, which showcases the top 10 topics for each dataset, ranked by the percentage of users who contributed at least one tweet on the topic.

Our analysis encompassed four datasets, comprising two polarized datasets (Gun and Abortion) and two unpolarized datasets (Super Bowl and SXSW). We observed that users within the polarized datasets exhibited a higher inclination towards political topics compared to the users in the unpolarized datasets. This observation confirms that discussions around political topics are more likely to exhibit characteristics of echo chambers [3].

Gun (ECS = 0.71)		**Abortion (0.63)**	
Politics*	99.3%	**Politics & Society***	99.0%
Social Media	76.1%	Social Media	91.4%
Celebrations	70.7%	Celebrations & Emotions	88.1%
Political Figures*	66.9%	Entertainment & Pets	86.8%
Entertainment	62.3%	**Health & Pandemic***	73.5%
Pandemic*	61.4%	Art & Lifestyle	71.0%
Lifestyle	57.1%	Space & Astrology	69.7%
Social Issues*	47.5%	Sports	46.9%
Weather	46.1%	**Energy & Climate***	38.2%
Pets & Wildlife	45.7%	Literature	34.9%
Sports	43.4%	Travel & E-commerce	25.5%

(a) Polarized Datasets

Super Bowl (0.48)		**SXSW (0.46)**	
Celebrations	79.2%	Celebrations	95.0%
Music	64.6%	Movies	77.0%
Communication	48.6%	Music	52.6%
Sports	36.4%	Cryptocurrency	43.6%
Football	34.2%	Social Media	40.3%
Social Media	33.4%	Social Issues	37.6%
Cryptocurrency	32.7%	Crypto & Art	28.0%
Entertainment	30.2%	**Elections***	26.4%
Social Issues*	21.4%	Technology	23.8%
Politics*	20.3%	**Public Policy***	12.7%
Royalty	19.4%	Green Energy	8.8%

(b) Unpolarized Datasets

Table 2: Top 10 topics for each dataset, along with their corresponding Echo Chamber Score (ECS) values. A higher ECS signifies increased polarization and a greater likelihood of the presence of an echo chamber.

Furthermore, our experiment suggests that users who engage with polarized topics are more likely to share and interact with other polarized topics, aligning with the concept of echo chambers. Although further experiments are necessary to establish a causal relationship between the interaction with polarized topics, it is reasonable to assert a correlation between the two.

Interestingly, in the unpolarized datasets, only a minority of users interacted with at least one polarized topic (21% in the Super Bowl dataset and 26% in SXSW). In contrast, nearly all users (99%) in the polarized datasets contributed at least one tweet on polarized topics such as politics.

Figure 1 provides insight into the interaction patterns, demonstrating that topics within the polarized datasets elicit greater engagement from users compared to the unpolarized datasets. Put simply, users within polarized datasets frequently interact with topics similar to those discussed by other users within the same dataset, indicating a narrower range of interests. This observation aligns with the characteristics of echo chambers and further suggests that polarization leads to a more limited interest in topics and a reduced diversity of discussions.

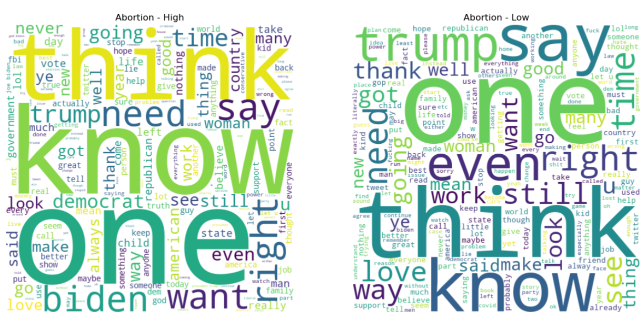

Fig. 2: Word Cloud for the two communities in Abortion Dataset (Abortion-High and Abortion-Low). We can see that just looking in the most frequent words tells very little about the topics of each community in comparison to the topics modeling based method (See Table 1).

3.2 Impact of Polarization on Community-Level Topic Distribution

This experiment focuses on studying the distribution of topics at the community level, utilizing the Echo Chamber Score (ECS) to measure polarization. Unlike other echo chamber methods that assess polarization at the graph level only, ECS enables us to measure polarization both at the graph level and the community level [2].

In this experiment, we specifically examine the topic distribution within two datasets: the Abortion dataset (polarized) and the SXSW dataset (unpolarized). By selecting one dataset of each type, we aim to determine whether the level of polarization of the dataset influences the topic analysis on the community level.

To detect communities within the datasets, we employ the widely used Louvain algorithm. Subsequently, we identify the most polarized and least polarized communities within each dataset. In the Abortion dataset, we observe two communities with sizes of 1,164 and 3,909 users, respectively, denoted as Abortion-High (based on the ECS score) and Abortion-Low. The ECS scores for these communities are 0.65 and 0.57, respectively. Similarly, within the SXSW dataset, we identify two large communities (with sizes of 1,606 and 677) named SXSW-High and SXSW-Low, with ECS scores of 0.488 and 0.45, respectively.

Abortion - High ECS = 0.65	Abortion - Low 0.57
Politics (95.1%)	**Politics (99.3%)**
Abortion (91.7%)	Greetings (91.9%)
Religion (90.8%)	Media (88.2%)
Ukraine War (90.4%)	Pets (79.8%)
Migration (79.8%)	Finance (79.4%)
Finance (76.7%)	**Covid 19 (71.8%)**
Covid 19 (68.7%)	**Climate Change (68.7%)**
Twitter (65.3%)	Fashion & Art (52.8%)
Pets (62.9%)	Entertainment (51.2%)
Bots (62.1%)	Sports (44.8%)

(a) Abortion dataset

SXSW - High ECS ECS = 0.49	SXSW - Low 0.45
Festivities (92%)	Community (93.3%)
Cinema (80.5%)	Celebrations (92.5%)
Social Media (77.8%)	NFT (70.7%)
Music (51.8%)	Sports (69.3%)
Social Issues (48.7%)	Cryptocurrency (60.6%)
Football (44.2%)	Trading (57.5%)
Fashion (41%)	Food (55.1%)
Communication (36.1%)	**Social Media (53.1%)**
Literature (34.4%)	Entertainment (49.7%)
Cryptocurrency (26.9%)	Life (44.5%)

(b) SXSW dataset

Table 3: The top 10 topics discussed in each dataset. We examine the communities with the highest and lowest ECS scores. For each dataset, we identify the community with the highest and lowest ECS score and identify the top 10 topics

We use BERTopic to examine the topics discussed within these communities. In contrast to the word clouds (Figure 2), BERTopic provides more granular insights into the topics of interest. In both the Abortion and SXSW datasets, we find that the more polarized communities (Abortion-High and SXSW-High) engage in more political discussions and exhibit greater topic homogeneity compared to the less polarized communities (Abortion-Low and SXSW-Low). Table 3 highlights this observation, with the polarized communities featuring a higher number of polarizing topics.

In conclusion, this experiment sheds light on the topic distribution at the community level, employing the ECS measure. The results indicate that polarization significantly influences the topic composition within communities, with more polarized communities exhibiting a greater tendency toward discussing specific topics, particularly in the realm of politics.

4 Conclusion and Future Work

In this study, our aim was to investigate the impact of polarization on the topics discussed within online communities. Through our analysis of dataset-level and community-level topic distributions, we have gained valuable insights into the prevalence and effects of echo chambers. Our findings highlight that individuals engaged in polarized discussions exhibit a stronger tendency to prioritize political topics, indicating the association between political subjects and the formation of echo chambers. Moreover, we observed greater topic homogeneity within more polarized communities, suggesting a higher susceptibility to the influence of echo chambers among users who share these topics. This study represents a pioneering attempt to examine the content shared by users within their timelines, shedding light on the dynamics surrounding echo chambers. However, further research is warranted to explore how users interact with these topics, analyze content and interaction patterns, and identify susceptible users and the emergence of potential echo chambers. Such investigations would contribute to a deeper understanding of echo chamber dynamics and facilitate the development of strategies to foster a more inclusive and informed public discourse. By addressing the harmful effects of echo chambers and promoting a diverse and balanced online environment, we can strive towards a healthier digital society.

Acknowledgement

This work was supported by the Saudi Arabian Cultural Mission and King Saud University. Also, this work was supported by the Office of Naval Research (ONR) under contract/grant number N00014-21-1-4002. Opinions, interpretations, conclusions, and recommendations are those of the authors.

References

1. Alatawi, F., Cheng, L., Tahir, A., Karami, M., Jiang, B., Black, T., Liu, H.: A survey on echo chambers on social media: Description, detection and mitigation (Dec 2021). https://doi.org/10.48550/arXiv.2112.05084

2. Alatawi, F., Sheth, P., Liu, H.: Quantifying the echo chamber effect: An embedding distance-based approach (Jul 2023). https://doi.org/10.48550/arXiv.2307.04668
3. Barberá, P., Jost, J.T., Nagler, J., Tucker, J.A., Bonneau, R.: Tweeting from left to right: Is online political communication more than an echo chamber? Psychological Science **26**(10), 1531–1542 (Oct 2015). https://doi.org/10.1177/0956797615594620
4. Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. Journal of Statistical Mechanics: Theory and Experiment **2008**(10), P10008 (Oct 2008). https://doi.org/10.1088/1742-5468/2008/10/P10008
5. Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., Starnini, M.: The echo chamber effect on social media. Proceedings of the National Academy of Sciences **118**(9), e2023301118 (Mar 2021). https://doi.org/10.1073/pnas.2023301118
6. Cota, W., Ferreira, S.C., Pastor-Satorras, R., Starnini, M.: Quantifying echo chamber effects in information spreading over political communication networks. EPJ Data Science **8**(1), 1–13 (Dec 2019). https://doi.org/10.1140/epjds/s13688-019-0213-9
7. Dahal, B., Kumar, S.A.P., Li, Z.: Topic modeling and sentiment analysis of global climate change tweets. Social Network Analysis and Mining **9**(1), 24 (Jun 2019). https://doi.org/10.1007/s13278-019-0568-8
8. Grootendorst, M.: Bertopic: Neural topic modeling with a class-based tf-idf procedure (Mar 2022). https://doi.org/10.48550/arXiv.2203.05794
9. Jiang, B., Karami, M., Cheng, L., Black, T., Liu, H.: Mechanisms and attributes of echo chambers in social media (Aug 2021). https://doi.org/10.48550/arXiv.2106.05401
10. Lin, C., He, Y.: Joint sentiment/topic model for sentiment analysis. In: Proceedings of the 18th ACM Conference on Information and Knowledge Management. pp. 375–384. CIKM '09, Association for Computing Machinery, New York, NY, USA (2009). https://doi.org/10.1145/1645953.1646003
11. OpenAI: Gpt-4 technical report (Mar 2023). https://doi.org/10.48550/arXiv.2303.08774
12. Salehi, A., Ozer, M., Davulcu, H.: Sentiment-driven community profiling and detection on social media. In: Proceedings of the 29th on Hypertext and Social Media. pp. 229–237. HT '18, Association for Computing Machinery, New York, NY, USA (Jul 2018). https://doi.org/10.1145/3209542.3209565