BeECD: Belief-aware Echo Chamber Detection over Twitter Stream

Guan Wang, Weihua Li, Shiqing Wu, Quan Bai, and Edmund M-K Lai

[1] Auckland University of Technology, Auckland, New Zealand
guan.wang@autuni.ac.nz{weihua.li, edmund.lai}@aut.ac.nz
[2] University of Tasmania, Hobart, Australia
quan.bai@utas.edu.au
[3] University of Technology Sydney, Sydney, Australia
shiqing.wu@uts.edu.au

Abstract. The phenomenon known as the "echo chamber" has been widely acknowledged as a significant force affecting society. This has been particularly evident during the Covid-19 pandemic, wherein the echo chamber effect has significantly influenced public responses. Therefore, detecting echo chambers and mitigating their adverse impacts has become crucial to facilitate a more diverse exchange of ideas, fostering a more understanding and empathetic society. In response, we use deep learning methodologies to model each user's beliefs based on their historical message contents and behaviours. As such, we propose a novel, content-based framework built on the foundation of weighted beliefs. This framework is capable of detecting potential echo chambers by creating user belief graphs, utilizing their historical messages and behaviours. To demonstrate the practicality of this approach, we conducted experiments using the Twitter dataset on Covid-19. These experiments illustrate the potential for individuals to be isolated within echo chambers. Furthermore, our in-depth analysis of the results reveals patterns of echo chamber evolution and highlights the importance of weighted relations. Understanding these patterns can be instrumental in the development of tools and strategies to combat misinformation, encourage the sharing of diverse perspectives, and enhance the collective well-being and social good of our digital society.

Keywords: Echo Chamber · Covid-19 · Belief Graph · Ego Network

1 Introduction

Nowadays, online social platforms have become one of the key sources for people to perceive information. It also reshapes the way of searching, filtering and disseminating information [17]. Modelling and analysing the influence and dissemination of information on social networks, including information maximization, social sentiment analysis, concern detection, etc. [12, 19], have become prominent subjects [11, 14]. One of the main perspectives of social media is to expose users to like-minded peers, which may result in echo chambers that could reinforce users' pre-existing viewpoints and drive the communities to be more polarized [7]. Individuals in these communities are easily affected by their surroundings. Estimating the extent to which an individual is isolated in an echo chamber is helpful in breaking the isolation.

Jamieson et al. were the first to define the phenomenon occurring on social media platforms where information within a community is amplified and metaphorically term it as the "echo chamber" [10]. To determine whether an individual is isolated within an echo chamber, we focus on their surroundings and behaviours. Therefore, we define the echo chamber from an individual perspective. This phenomenon is characterized by an individual who:

- Resides in a community that echoes their opinion, wherein most neighbours share similar views.
- Inhabits this community, where the individual's perspectives are repeatedly reinforced by community members.
- Self-reinforces by engaging in communication that aligns with their viewpoint.

The prominence of the echo chamber phenomenon has been heightened by the outbreak of Covid-19. Recognised globally as a pandemic, Covid-19 has become a leading topic of discussion. Conversations around this subject vary widely, encompassing themes such as vaccine hesitancy and vaccination-related deaths. Under these conditions, social platforms provide a conducive environment for misinformation propagation due to the lack of editorial supervision. As a result, echo chambers have emerged among users, significantly influencing responses to the Covid-19 pandemic [2]. For example, if members of a community consistently engage with and promote content sceptical of Covid-19 vaccinations, the community can be highly identified as an echo chamber resistant to the prevailing medical advice on vaccines.

To address this challenging issue, we propose a novel approach to detect the echo chamber phenomenon and estimate the corresponding degree. In pursuing this aim, we explore the utility of knowledge graphs in identifying echo chambers from an individual perspective and propose a content-based framework that constructs belief graphs for each individual. During this construction process, we extract triplets from the individual's related Twitter content, including tweets, retweets, and replies. To assess the impact of individual behaviours, we incorporate them as distinct parameters to calculate weights for these triplets. Extensive experiments are conducted, and the results explicitly show that modelling an individual's belief graphs with weighted relations can effectively reveal an individual's trends on a specific topic and identify the echo chambers to which the individual belongs.

2 Related Work

2.1 Echo Chambers on Social Platforms

Social platforms have subtly altered the way people access information and formed echo chambers, isolating them in the process [7]. The detection of echo chambers has been a research focus across various fields [4, 18], serving as the first step towards mitigating this phenomenon.

Social structures typically manifest in two distinct forms, i.e., global perspective and individual perspective. Most research has studied echo chambers from a global perspective or topological viewpoint, primarily focusing on user interactions while overlooking the source of these interactions [5, 6]. Cinelli et al. analyze echo chambers by assessing whether the overall network is strongly polarized towards two sides of a controversy, emphasizing user interaction networks [7]. Cossard et al. explore echo chambers within vaccine communities using clustering techniques, demonstrating the existence of echo chambers within real social networks [8].

Analyzing extensive topological structure datasets from a global perspective necessitates high-performance computing resources. The ego network centred around a focal user offers a feasible way to model a community, enabling measurement of the echo chamber degree with a focus on that user. Thus, inspired by Li et al. and Valerio et al., we incorporate the concept of the ego network in our study [3, 13]. Li et al. propose agent-based influence diffusion models, where the influence cascading process is modelled as an evolutionary pattern driven by individuals' actions. Valerio et al. analyze the micro-level structural properties of online social networks and demonstrate that ego networks play a significant role in social networks, impacting information diffusion within the network.

2.2 Content-Based Echo Chamber Detection

Content-based methods identify echo chambers by analyzing the information texts produced by individuals. Villa et al. propose both a topology-based and content-based approach, analyzing the topological structure of the social network and sentiment aspects related to the content [21]. Cinelli et al. conduct a comparative analysis on a large-scale dataset to identify echo chambers through social network homophily. They define "leaning" as the attitude expressed by a piece of content towards a specific topic about the content [7]. Abd-Alrazaq et al. propose a text-mining method on a large dataset, considering information texts but neglecting temporal information, which can provide contextual insights [1]. Lwin et al. and Xue et al. demonstrate that discourses on Twitter about Covid-19 continually evolve, develop, or change over time [15, 22]. Inspired by these studies, we restructure the dataset into chronologically user-specific streams.

Most existing studies solely consider the content of information but overlook individual behaviours and content weights, which demonstrate the significance of content on individuals. For instance, reading a message doesn't explicitly reveal an individual's thoughts about the message. However, a subsequent 'like' or 'upvote' implies that the individual agrees with this message, thereby increasing the weight of information from this message in the corresponding belief graph. We argue that beliefs in individuals' minds carry different weights, and not all beliefs hold equal significance. As a result, behaviours offer valuable insights into people's perspectives on related messages.

Therefore, we propose a belief-aware echo chamber detection framework incorporating content and individual behaviours. Our framework constructs belief graphs for each individual in our dataset, considering their behaviours. To measure the degree of echo chambers, we calculate the similarities between the belief graphs of the focal user and their neighbours. With this framework, social platforms can detect communities where members are primarily exposed to reinforcing views, potentially limiting the diversity of thoughts and contributing to polarization.

3 Belief-Aware Echo Chamber Detection

In this section, we formally define related terms and explain the proposed belief-aware echo chamber detection framework.

3.1 Formal Definitions

Two types of graph structures are utilized in this work: one is the ego network, a directed graph $G =< U, E >$ that includes a focal user and its neighbours, and the other is the user belief graph $BG =< H, WR, T >$.

Definition 1. ***An ego network*** *consists of a focal user* u_f *and their neighboring users, denoted as* $U = \{u_0, ..., u_n\}$*. Each user, represented by* $u_i \in U$*, corresponds to a node in this directed network. Each edge*(u_i, u_j) *between nodes is directed as the flow of information, indicating that user* u_i *follows, replies to, or mentions user* u_j*. Each user in an ego network also has a unique belief graph, representing their personal network of beliefs, which is clarified in Definition 2.*

Definition 2. ***A belief graph*** *is a unique graph containing multiple triplets* $BG = \{H, WR, T\}$*, where* H *and* T *refer to nodes in belief graphs, and* $WR = \{wr_i | 0 < i < m\}$ *represents relations between nodes, defined as weighted relations. The belief graph is constructed by extracting triplets from users' historical messages and behaviours on corresponding messages across various topics.*

Definition 3. ***Similarity*** $sim(v_i, v_j)$ *refers to the distance between two vectors in a low-dimensional space. The similarity is a value between* $[0, 1]$*, where* 0 *implies completely contrary viewpoints, while* 1 *signifies identical viewpoints.*

Definition 4. ***Echo chamber degree*** $p(u, k)$ *is a measure that evaluates the likelihood of an echo chamber. A higher degree suggests a higher probability that an individual experiences an echo chamber related to a specific topic* k*.*

Definition 5. ***Topic*** k *refers to the label of each message, e.g.,* m_k*. The topic set is denoted as* $T = \{T_0, T_1, ..., T_n\}$*. Messages with the same topics express similar discourse. One message is assigned only one topic. In our framework, topics are used for graph partitioning.*

As illustrated in Fig. 1, we determine the echo chamber degree in 5 phases:

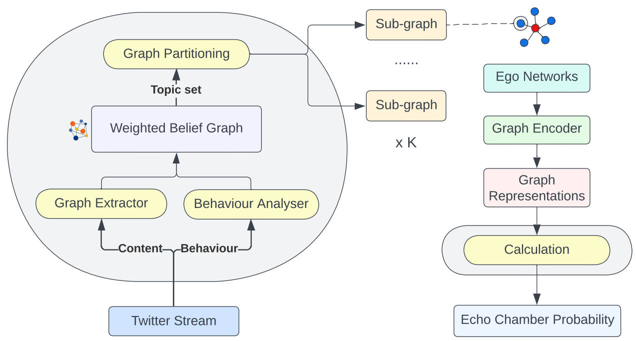

Fig. 1. The brief overall process of the framework.

- 1) Construct a belief graph for each individual, considering both their Twitter stream and their user behaviours.
- 2) Partition the corresponding part of an individual's belief graph into sub-graphs according to different topics.
- 3) Select an individual and its neighbours to create an ego network based on their followee/follower relations and mentioning behaviour.
- 4) Calculate the similarities of sub-graphs on a given topic between the focal user and its neighbours to assess the closeness of their beliefs.
- 5) Quantify the echo chamber degree by evaluating their average similarity and information entropy.

A sub-graph is a graph partitioned from an individual's complete belief graph given a specific topic. Sub-graphs are used to compare users. Messages refer to texts that users receive and post, including tweets, retweets, and replies. User behaviours refer to user operations on a social platform, including:

- Viewing: Users view messages posted by their neighbours (followees) or recommended by the platform.
- Liking: Users like a message by clicking the blank heart symbol.
- Disliking: Users express dislike for a message by cancelling their liking behaviour, i.e., clicking the solid heart symbol.
- Reposting: Users repost viewed messages.
- Sending: Users post a message or reply to someone in their own words.

We presume each behaviour reflects a different perspective on corresponding messages and aids in modelling changes in user beliefs. For example, when a user likes a message, we increase the weight of triplets extracted from this message by assigning a changing rate to the weight. The changing rates are defined for different behaviours as shown in Table 1:

3.2 Belief Graph Construction

The first step in this work is to construct belief graphs for each individual in the ego network. We extract triplets (i.e., $\{$*head, relation, tail*$\}$) from the content

Table 1. Changing rate of each behaviour on corresponding information.

	Reviewing	**Liking**	**Disliking**	**Reposting**	**Sending**
Changing Rate(r)	0.5	2	-2	1	2

and calculate weights for these triplets by analyzing the individual's behaviours. We then attach the weights to the relations, resulting in weighted relations. A belief graph that reflects an individual's belief consists of multiple triplets with weighted relations. We use Stanford OpenIE [1] to extract triplets from texts.

To calculate the weights, we employ a logarithmic function to prevent the weights from reaching extremely high or low. This logarithmic transformation helps maintain a balanced range of weights. The function is defined as follows:

$$w_i = ln(w_i^{'} + r_i) + 1, \quad (1)$$

where $w_i^{'}$ is the previous weight of the same triplet and r_i is the changing rate as shown in Table 1.

During extraction, the same triplets may be extracted multiple times. For each new triplet, the initial weight is defined as 0, and its current weight is calculated based on the changing rate of the corresponding behaviour. When we encounter the identical triplet, we add the change in weight according to the current behaviour to its previous weight (i.e., $w_i^{'}$ in Equation 1).

3.3 Belief Graph Partitioning

A complete belief graph of an individual encompasses information from several topics. Comparing complete graphs may allow irrelevant information to affect performance on the given topic. Hence, we perform a graph partitioning step before transforming graphs into graph representations.

In this step, we utilize word embeddings from a word2vec model [16] to identify nodes within the belief graph that have similar words to the given topic and keywords. The cosine similarity function, as shown below, is utilized to measure the similarity between word embeddings. Nodes and relations in both directions are subsequently used to form the sub-graph.

$$sim(v_i, v_j) = \frac{v_i \cdot v_j}{\parallel v_i \parallel \parallel v_j \parallel}, \quad (2)$$

where v_i and v_j denote word embeddings obtained from the word2vec model.

3.4 Echo Chamber Detection

To compare the similarities among belief graphs, we convert these topological structure graphs into vector representations. This is achieved through training Graph Attention Networks (GATs) on each individual's belief graph to generate

[1] https://nlp.stanford.edu/software/openie.html

graph representations. Different from the original GATs [20], we introduce the weighted relation features $R = \{r_{i,j} | 0 < i < n, 0 < j < n\}$ as the initial attention coefficient. The weighted relation features are used during the attention calculation as follows:

$$e_{i,j} = a(W\hat{h}_i, W\hat{h}_j, r_{i,j}) \tag{3}$$

Equation 3 represents the importance of node j's features to node i. W denotes a weight matrix used to parameterize a shared linear transformation, $\hat{h}_i$ represents the features of node i, and $r_{i,j}$ is the weighted relation from node i to node j. To collect all features of the whole graph, we add a global node to each graph. This global node is linked to every node in the graph, and its representation represents the entire graph.

We hypothesize that similar graphs express similar beliefs on relevant topics. To test this, we compute the similarity between individuals' belief graphs. We apply the graph representations generated by the trained GATs to a cosine similarity function to calculate these similarities:

$$sim(h_i^k, h_u^k) = \frac{h_i^k \cdot h_u^k}{\| h_i^k \| \| h_u^k \|}, \tag{4}$$

where h_i^k denotes the representation of user i's sub-belief graph on topic k, and $\| h_i^k \|$ represents the Euclidean norm of h_i^k. h_u^k refers to the representation of the focal user's sub-graph on topic k. The average similarities are then calculated as follows:

$$avg(h_u^k) = 1/n \sum_{i=1}^{n} sim(h_i^k, h_u^k), \tag{5}$$

where n describes the number of the focal user's neighbours.

In addition to similarity, inspired by [9], we also consider information entropy from information theory and statistical mechanics to calculate the probability of an individual being isolated in an echo chamber. The equation is as follows:

$$H(g_u^k) = -\sum_{k \subseteq K} p_k \cdot ln(p_k), \tag{6}$$

where p_k is the percentage of a user's sub-graph on topic k, and g_u^k refers to the belief graph of the focal user of an ego network on topic k. Finally, we use both average similarity and information entropy to measure the echo chamber using the following equation:

$$p(u,k) = avg(h_u^k) \cdot H(g_u^k), \tag{7}$$

where u represents a focal user. A higher $p(u,k)$ indicates a greater likelihood that u is isolated in an echo chamber. In such a case, the ego network centred around user u is a $p(u,k)$ possibility echo chamber on topic k.

4 Experiments and Analysis

This section provides details of two experiments conducted to validate the efficacy of the proposed Belief-based Echo Chamber Detection model. The first experiment evaluates the similarities in responses of echo chamber members to multiple related messages. The second experiment implements an ablation study to elucidate the progression of the Belief Graph module within the BeECD framework and to investigate the impact of weighted relations on belief graphs.

4.1 Data Collection and Organisation

The experiments utilize a dataset gathered from Twitter related to Covid-19. This dataset, part of the continually updated Covid-19 Twitter chatter dataset maintained by Georgia State University's Panacea Lab, spans a crucial six-month period from December 2020 to May 2021. This time frame is particularly significant as it encompasses a period when several candidate vaccines displayed safety and the ability to generate immune responses. The proposed BeECD can be applied to any dataset. In this paper, we leverage Covid-19 as the dataset to validate this approach.

Each individual's content and behaviours are organised into a chronological stream, including the user's tweets, tweets from the user's neighbours, retweets, replies, corresponding tweets, likes, and liked tweets. We limit our focus solely to English tweets, replies, and retweets, and uniquely, we include retweets in the streams of each individual, allowing us to process retweeting behaviour and corresponding content concurrently.

To facilitate computation, we extracted a sub-graph from the total dataset, comprising 285 users, 3,587 interconnections, and 42,478 posts, which include tweets, retweets, and replies.

4.2 Experiment 1: Response Analysis

In this experiment, we hypothesize that each user within an ego network can respond appropriately to one or multiple similar messages, anticipating that responses from like-minded users will exhibit greater similarity than those from dissimilar users. The implications of this phenomenon in real-world contexts are substantial. Consider, for instance, an ego network exhibiting an 80% echo chamber probability. If the average response probabilities within this network align closely with this percentage, it will signify that most users within the network are engaged in disseminating and consuming similar information. In a practical sense, this may translate to a reinforcement of a specific narrative or perspective. The resulting lack of engagement with diverse viewpoints could amplify polarization. This may create a self-reinforcing cycle in which users are confined to information confirming their pre-existing beliefs, thereby becoming increasingly resistant to alternative viewpoints or evidence contradicting their established convictions.

In addition to calculating the echo chamber probabilities, we subsequently train an encoder-decoder structured language model on the entire dataset, feeding 20 random messages from the test set into each ego network based on the same topic that the ego network inclines towards. The language model is used to assess the similarity of the users' responses. By comparing the echo chamber probabilities and response similarities, we assess the efficacy of the proposed framework. The results depicted in Fig 2 corroborate our hypothesis.

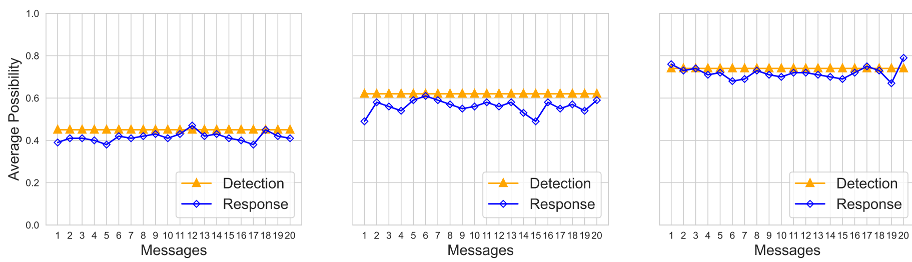

Fig. 2. The probabilities of detected echo chamber and user response similarities.

By presenting the outcomes from three distinct ego networks with varying degrees of echo chamber probabilities, it's clear that the average response probabilities align with the calculated echo chamber probabilities.

This experiment sheds light on the intricate relationship between online interactions and the formation of echo chambers. The observed alignment between the echo chamber probability and user response patterns underscores the role digital platforms play in shaping real-world perspectives, emphasizing the need for further research and interventions in this domain.

4.3 Experiment 2: Belief Graph Impact Analysis

The second experiment seeks to understand the influence of weighted relations on the evolution of echo chambers. We additionally train Graph Attention Networks (GATs) that do not account for the properties of relations during the training process. Belief graphs are initiated using data from the first two months of six, following which the remaining data is partitioned into 25 unique time intervals. From each interval of the users' Twitter streams, users' beliefs and behaviours are extracted and used to update their corresponding belief graphs.

By partitioning data into distinct time intervals, we highlight the significance of temporal evolution in shaping users' beliefs. This process shows the importance of identifying the beliefs and understanding how they transform and develop over time. Such an approach corresponds to real-world scenarios, where individuals often undergo phases or shifts in their perspectives. These changes may be influenced by various factors, such as past experiences, exposure to new information, or personal growth, reflecting the complexity and dynamism of human belief systems.

We anticipate that our framework, which incorporates weighted relations, is capable of detecting variations in user beliefs, including instances where these beliefs intensify before subsequently diminishing. This approach offers insight into the dynamic nature of belief changes. On the other hand, in the absence of such weighted relations, a user's beliefs appear to remain unaltered and static. This lack of dynamism obscures the potential to observe the evolutionary patterns of echo chambers, even in cases where users undergo significant shifts in their perspectives. We selected four representative curves from our proposed framework with weighted relations and corresponding curves from the framework lacking weighted relations for comparison. The results depicted in Fig 3 effectively showcase different patterns of evolution of echo chambers in our proposed framework with weighted relations.

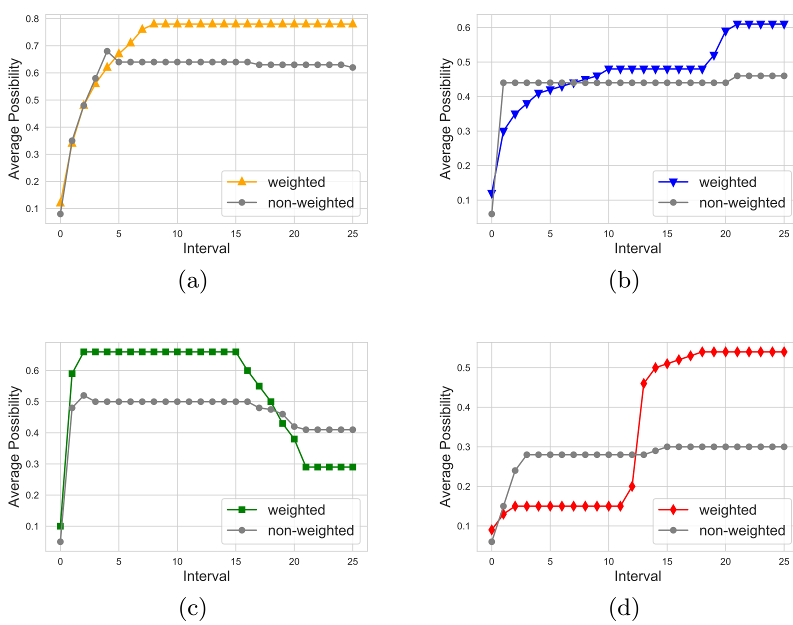

Fig. 3. The evolution of echo chambers.

From Fig 3, it is clear that our proposed framework can effectively represent the evolution of echo chambers or users' shifting perspectives. The result also highlights a crucial difference between the two models. The use of weighted relations, as opposed to non-weighted ones, allows for a more nuanced representation of the complexities inherent in human interactions and belief systems. In real-world terms, not all interactions influence our beliefs equally. Some might have a significant impact due to the trustworthiness of the source or the emotional resonance of the content, while others might be casually scrolled past without much thought. Thus, incorporating weighted relations can more accurately model how real people might be influenced by their digital interactions.

Understanding the evolution of echo chambers using advanced models like GATs with weighted relations is crucial in today's digital age. Such insights provide a clearer picture of how beliefs change over time on platforms like Twitter, emphasizing the need for digital platforms to prioritize diverse content exposure and critical thinking among their users.

5 Conclusion and Prospective Research Directions

In this study, we introduce a novel content-based methodology for echo chamber detection on social networks, coined as the belief-aware echo chamber detection approach shedding light on the intricate relationship between online interactions and the formation of echo chambers. We leverage Knowledge Graph technology to construct user belief graphs, taking into account both message content and user behaviour. Additionally, we train modified Graph Attention Networks, incorporating weighted relations into the computation process. Similarities between user belief graphs are then computed. The experimental results indicate promising effectiveness and demonstrate real-world implications of our approach in analyzing echo chambers on social platforms.

However, the detection of echo chambers represents a seminal work, and addressing the subsequent effects presents significant challenges. Future research endeavours will focus on strategies for mitigating the echo chambers, further advancing the understanding and management of social network dynamics.

References

1. Abd-Alrazaq, A., Alhuwail, D., Househ, M., Hamdi, M., Shah, Z., et al.: Top concerns of tweeters during the covid-19 pandemic: infoveillance study. Journal of medical Internet research **22**(4), e19016 (2020)
2. 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. arXiv preprint arXiv:2112.05084 (2021)
3. Arnaboldi, V., Conti, M., La Gala, M., Passarella, A., Pezzoni, F.: Ego network structure in online social networks and its impact on information diffusion. Computer Communications **76**, 26–41 (2016)
4. Bail, C.A., Argyle, L.P., Brown, T.W., Bumpus, J.P., Chen, H., Hunzaker, M.F., Lee, J., Mann, M., Merhout, F., Volfovsky, A.: Exposure to opposing views on social media can increase political polarization. Proceedings of the National Academy of Sciences **115**(37), 9216–9221 (2018)
5. 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 (2015)

6. Bruns, A.: Echo chamber? what echo chamber? reviewing the evidence. In: 6th Biennial Future of Journalism Conference (FOJ17) (2017)
7. Cinelli, M., Morales, G.D.F., Galeazzi, A., Quattrociocchi, W., Starnini, M.: The echo chamber effect on social media. Proceedings of the National Academy of Sciences **118**(9) (2021)
8. Cossard, A., Morales, G.D.F., Kalimeri, K., Mejova, Y., Paolotti, D., Starnini, M.: Falling into the echo chamber: the italian vaccination debate on twitter. In: Proceedings of the International AAAI conference on web and social media. vol. 14, pp. 130–140 (2020)
9. Hu, Y., Wu, S., Jiang, C., Li, W., Bai, Q., Roehrer, E.: Ai facilitated isolations? the impact of recommendation-based influence diffusion in human society. In: Raedt, L.D. (ed.) Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22. pp. 5080–5086. International Joint Conferences on Artificial Intelligence Organization (2022)
10. Jamieson, K.H., Cappella, J.N.: Echo chamber: Rush Limbaugh and the conservative media establishment. Oxford University Press (2008)
11. Jiang, C., D'Arienzo, A., Li, W., Wu, S., Bai, Q.: An operator-based approach for modeling influence diffusion in complex social networks. Journal of Social Computing **2**(2), 166–182 (2021)
12. Li, W., Bai, Q., Jiang, C., Zhang, M.: Stigmergy-based influence maximization in social networks. In: PRICAI 2016: Trends in Artificial Intelligence: 14th Pacific Rim International Conference on Artificial Intelligence, Phuket, Thailand, August 22-26, 2016, Proceedings 14. pp. 750–762. Springer (2016)
13. Li, W., Bai, Q., Zhang, M.: Agent-based influence propagation in social networks. In: 2016 IEEE International Conference on Agents (ICA). pp. 51–56. IEEE (2016)
14. Li, W., Bai, Q., Zhang, M.: Siminer: a stigmergy-based model for mining influential nodes in dynamic social networks. IEEE Transactions on Big Data **5**(2), 223–237 (2018)
15. Lwin, M.O., Lu, J., Sheldenkar, A., Schulz, P.J., Shin, W., Gupta, R., Yang, Y.: Global sentiments surrounding the covid-19 pandemic on twitter: analysis of twitter trends. JMIR public health and surveillance **6**(2), e19447 (2020)
16. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space (2013)
17. Morini, V., Pollacci, L., Rossetti, G.: Toward a standard approach for echo chamber detection: Reddit case study. Applied Sciences **11**(12), 5390 (2021)
18. Romer, D., Jamieson, K.H.: Patterns of media use, strength of belief in covid-19 conspiracy theories, and the prevention of covid-19 from march to july 2020 in the united states: Survey study. Journal of medical Internet research **23**(4), e25215 (2021)
19. Shi, J., Li, W., Yang, Y., Yao, N., Bai, Q., Yongchareon, S., Yu, J.: Automated concern exploration in pandemic situations-covid-19 as a use case. In: Knowledge Management and Acquisition for Intelligent Systems: 17th Pacific Rim Knowledge Acquisition Workshop, PKAW 2020, Yokohama, Japan, January 7–8, 2021, Proceedings 17. pp. 178–185. Springer (2021)
20. Veličković, P., Casanova, A., Liò, P., Cucurull, G., Romero, A., Bengio, Y.: Graph attention networks (2018)
21. Villa, G., Pasi, G., Viviani, M.: Echo chamber detection and analysis. Social Network Analysis and Mining **11**(1), 1–17 (2021)
22. Xue, J., Chen, J., Chen, C., Zheng, C., Li, S., Zhu, T.: Public discourse and sentiment during the covid 19 pandemic: Using latent dirichlet allocation for topic modeling on twitter. PloS one **15**(9), e0239441 (2020)