

Article

Not peer-reviewed version

---

# How to find Orchestrated Trolls?A Case Study on Identifying Polarized Twitter Echo Chambers

---

[Nane Kratzke](#) \*

Posted Date: 2 February 2023

doi: 10.20944/preprints202302.0032.v1

Keywords: Social network; Twitter; Structural analysis; Echo chamber; Detection; Case study; German language; Disinformation



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

## Article

# How to Find Orchestrated Trolls? A Case Study on Identifying Polarized Twitter Echo Chambers

Nane Kratzke 

Lübeck University of Applied Sciences, Department for Electrical Engineering and Computer Science,  
Mönkhofer Weg 239, 23562 Lübeck, Germany

\* Correspondence: nane.kratzke@th-luebeck.de

**Abstract:** **Background:** This study presents a graph-based and purely structural analysis to detect echo chambers on Twitter. Echo chambers are a concern as they can spread misinformation and reinforce harmful stereotypes and biases in social networks. **Methods:** The study recorded the German-language Twitter stream over two months, recording about 180.000 accounts and their interactions. The study focuses on retweet interaction patterns in the German-speaking Twitter stream and found that the greedy modularity maximization and HITS metric are the most effective methods for identifying echo chambers. **Results:** The purely structural detection approach was able to identify an echo chamber (red community) that was focused on a few topics with a triad of Anti-Covid, right-wing populism, and pro-Russian positions (very likely reinforced by Kremlin-orchestrated troll accounts). In contrast, a blue community was much more heterogeneous and showed "normal" communication interaction patterns. **Conclusions:** The study highlights the effects of echo chambers as they can make political discourse dysfunctional and foster polarization in open societies. The presented results contribute to identifying problematic interaction patterns in social networks often involved in the spread of disinformation by problematic actors. It is important to note that not the content but only the interaction patterns would be used as a decision criterion, thus avoiding problematic content censorship.

**Keywords:** social network; Twitter; structural analysis; echo chamber; detection; case study; German language; disinformation

## 1. Introduction

Echo chamber detection in social networks refers to identifying and analyzing the formation and reinforcement of homogeneous groups in online communities, where individuals with similar beliefs and attitudes tend to congregate and reinforce each other's views. This phenomenon is a concern as it can lead to the spread of misinformation and the reinforcement of harmful stereotypes and biases. Echo chambers are particularly likely in areas where right-wing populist content is disseminated via social media [1,2]. The awareness of echo chambers and their effects has increasing importance in democratic and free societies, as echo chambers can partly make the necessary political discourse dysfunctional [3,4] and can foster polarization [5,6] in open societies. Many studies indicate that echo chambers in social networks might affect democratic elections [7–9]. Particularly in the case of very close election outcomes, such as the BREXIT [10–12], these communication patterns can be decisive and exploited by non-democratically legitimized actors (even foreign powers with manipulative tendencies) to pursue their interests, which benefit from divided societies or exploit them for their purposes [13]. Kremlin-orchestrated troll accounts, in particular, have gained sad notoriety here, exploiting these mechanisms to shape opinion in democratic societies [14–16].

Thus, there are significant reasons for free societies to understand and manage these mechanisms so that free speech can continue to unfold its constructive power and not be abused to hatch "manipulative cuckoo eggs" of foreign powers or non-democratically legitimized actors. To this end, it is essential to understand, among other things, that disinformation works particularly well in echo chambers. To recognize echo chambers is, therefore, of considerable interest. However, there

are several challenges in detecting echo chambers in social networks. One major challenge is the lack of clear definitions and criteria for identifying echo chambers and the difficulty in quantifying the extent to which an online community is homogeneous. Additionally, echo chambers can occur at different granular levels, such as individual users, groups, or entire platforms. The state-of-the-art methods for echo chamber detection in social networks involve using computational techniques to analyze the structure and content of online communities, such as network analysis, natural language processing, and machine learning. These methods have been applied to various platforms, such as Twitter, Facebook, and Reddit [5,17,18].

While content-based identification is more reliable, it is also more laborious and dependent on a large amount of labelled training data which often does not exist or has to be recorded and labelled elaborately. So, content-considering methods work very well for past events with sufficient data but have problems adapting to new circumstances (such as the Russia-Ukraine war). Especially the automated detection of orchestrated troll accounts is difficult in the details and usually does not lead to reliable results [19–22].

This paper, therefore, investigates whether purely structural analysis of communication and interaction patterns can be used to identify echo chambers. Because social networks like Twitter provide samples of their data streams for free, the study was also interested in whether this could be done with a sample of network interactions. E.g., the public Twitter event stream API provides a 1%-sample of the actual event data stream. And studies like [23] show that the provided samples are representative.

**Research Question 1:** *Can echo chamber detection be accurately performed using graph-based analysis of a sample of network interactions without the need for additional information, such as the content of the interactions or characteristics of the individuals involved?*

The object of investigation was the public sample of events of Twitter retweets of the German-speaking language room. This study thus recorded the public German Twitter datastream of interaction events over several months. Users can post short messages on Twitter, called "tweets," which can include text, images, and links. These tweets are displayed on the user's profile and can be seen by their followers.

- A **"status post"** is a tweet that a user composes and shares on their own profile. It can include text, images, and links and can be seen by their followers. Status posts are the starting point of information distribution on social networks like Twitter. Their content can only be captured and processed using complex Natural Language Processing (NLP) methods.
- A **"reply"** is a tweet a user composes in response to another user's tweet. When a user replies to a tweet, the original tweet is linked within the reply so that others can see the context of the conversation. Replies can also be seen by the followers of the user who wrote the original tweet. Replies can be confirming, questioning, contradicting, referring, and of course, any other form. Consequently, these interactions also require complex NLP methods to classify the interaction's character.
- A **"retweet"** is when a user shares another user's tweet on their profile. Retweets allow users to share content from other users with their followers. The analytical advantage of retweets is that content is shared without additional remarks or annotations. So one can assume that the retweeter agrees with the original tweet's content. Due to the accumulation of retweet interactions between the same accounts, it can be assumed that the content of these accounts is close to each other without having to analyze the actual content.
- A **"quote"** is similar to a retweet, but instead of simply sharing the original tweet, the user includes it as a quote in their tweet, along with their commentary. This allows users to share and comment on tweets in a way that allows the context of the original tweet to remain visible. Unlike a retweet, the original content is accompanied by comments that can change the meaning of the

original tweet from case to case. This possible change in meaning can be sarcasm, corrections, annotations, etc., which usually require complex content-based analysis using NLP methods.

These interactions allow users to share content and engage in conversations with other users on the platform. Especially retweets are a powerful tool for identifying communities on Twitter because they indicate that a user is interested in and endorsing the content of another user's tweet. When users retweet a tweet, they essentially share it with their followers and endorse it as something they find valuable or interesting. Analyzing retweeting patterns among users makes it possible to identify groups of users with similar interests and share similar content. Retweets can also be used to identify the most influential members of a community. Users frequently retweeted by others are likely to be seen as leaders or experts within a community, and their tweets are likely to be more widely seen and shared. Identifying these influential users is possibly better for understanding a particular community's dynamics and interests.

Additionally, retweets can be used to trace the spread of information and ideas within a community. When a tweet is retweeted, it is exposed to a new group of users who may be interested in the same topic. By analyzing patterns of retweets, it is possible to see how information and ideas spread within a community and how different groups of users influence the conversation. In summary, focusing on retweets is a viable approach to detecting communities on Twitter because it allows the identification of groups of users with similar interests, the detection of influential members, and traces the spread of information and ideas within a community. Therefore, this study focuses on retweet interactions because of the analysis's simplicity. However, this leads to a second research question:

**Research Question 2:** *How much of the available information is used for echo chamber detection if only retweet interactions from a representative sample are taken?*

## 2. Materials and Methods

Structural analysis of echo chambers involves analyzing the network structure of online communities to identify patterns of homogeneity and homophily (the tendency for individuals to associate with similar others). This typically includes techniques like community detection, centrality measures, and network visualization.

- One popular method for community detection is the use of modularity optimization algorithms [24], which aim to identify groups of nodes (representing individuals or groups) that are more densely connected than the rest of the network. This can reveal the existence of echo chambers within a larger online community.
- Centrality [25,26] measures can also be used to identify critical actors within an echo chamber, such as those with high betweenness centrality (i.e. individuals who bridge multiple groups) or those with high eigenvector centrality (i.e. individuals who are connected to highly central actors).
- Network visualization techniques [27], such as graph layouts [28] and node colouring, can also be used to reveal patterns in the structure of echo chambers (see Figure 6).

The benefits of using structural analysis to detect echo chambers include identifying homogeneous groups within a larger community and identifying critical actors within those groups. Additionally, network analysis can be applied to various platforms and at different levels of granularity. However, there are also drawbacks to this approach. One limitation is that structural analysis alone may not provide insight into the content or dynamics of echo chambers, so other methods, such as natural language processing or sentiment analysis, may also be needed. The choice of community detection algorithm or centrality measure can also affect the results obtained.

To systematically get a handle on all these degrees of analytical freedom, the process shown in Figure 1 was followed.

- Sec. 2.1 explains how the Twitter event stream has been recorded for follow-up in-depth analysis.
- Sec. 2.2 shows how the recorded Twitter interactions were analyzed. Among other things, Sec. 2.2.1 shows how a retweet graph can be constructed from a stream of Twitter interactions and how the largest components within this graph are determined to focus on the relevant part of the analysis. Sec. 2.2.2 focus on determining communities within the largest component of the graph. And finally, Sec. 2.2.3 shows how the different components can be measured to determine a characteristic property of echo chambers.
- Finally, Sec. 2.3 explains the evaluation of the presented approach using consecutive timeframes. The corresponding results are shown in Sec. 3.




Figure 1. The methodology followed in this study.

2.1. Recording

This study recorded a Twitter event stream (see Figure 2) using the Public Twitter Streaming API using the Python library Tweepy [29], and the document store database MongoDB [30]. This approach has been successfully used in previous studies [31] of the author.




Figure 2. Archtitecture used to record German-language Twitter traffic via the Public Streaming API. The Public Streaming API provides a representative sample of about 1% of the total traffic.

To repeat the experiments, one has to set up (or use an existing) Twitter account with the Public Twitter Streaming API. This requires requesting a developer account and obtaining an API key, which can be used to access the API. For the recording, this study used the Python Tweepy library. Tweepy is a Python wrapper for the Twitter API that makes it easy to access and manipulate data from the API. A Python script uses the Tweepy library to access the Public Twitter Streaming API and collect data from the event stream (see Listing 1 in Appendix A). The script should include the API key and other necessary parameters that will be used to filter the event stream. Once the script runs, it will continuously collect data from the event stream and store it in a JSON-based document store database. This study used MongoDB, a popular JSON-based document store database that can store and query large amounts of data. It is a NoSQL database and can be easily integrated with Python. The data stored in MongoDB can be queried and analyzed using various analytical tools such as Pandas [32], Spark [33] and many more.

In this study, tweets were recorded in a Kubernetes cluster over several months. The recorder and the database were run in containers. Kubernetes [34] ensured that the corresponding components were automatically restarted if the recorder or the database failed. However, the components can also be operated non-containerised in other environments. The following resources have been assigned to the recorder and the database.


- Recorder: 250m CPU, 250MB Memory
- MongoDB: 8000m CPU, 60GB Memory



2.2. Analysis

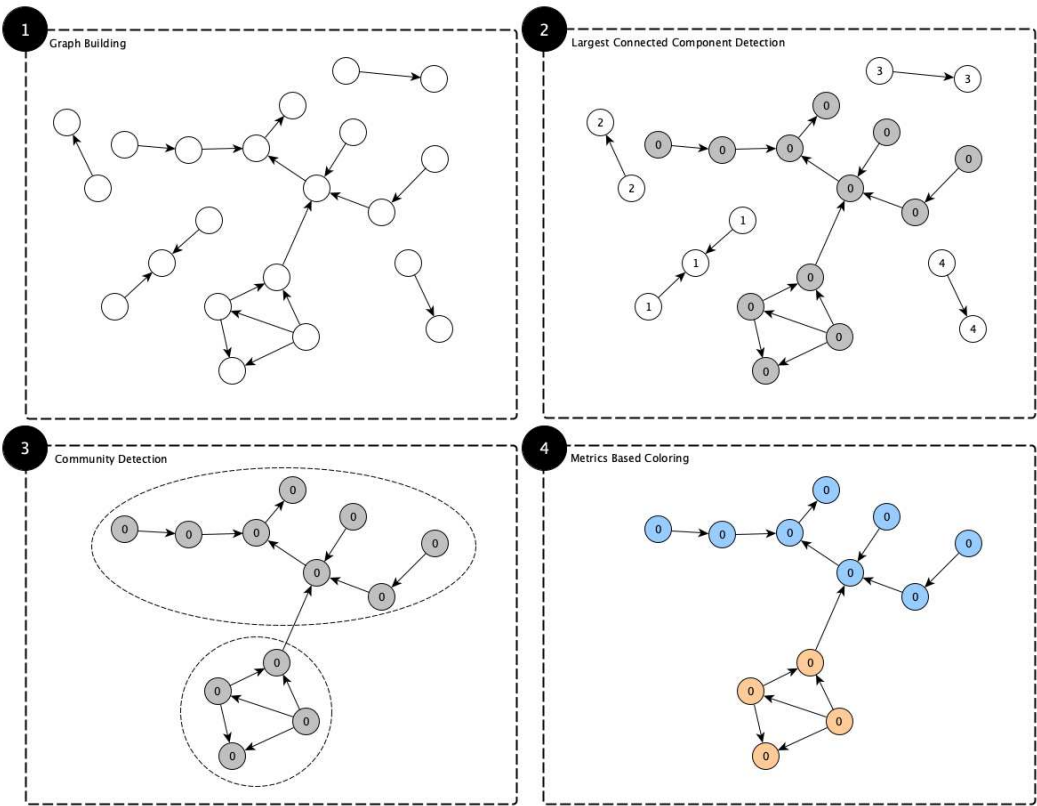
The basic workflow of analyzing Twitter interactions followed the procedure shown in Figure 3.

1. In the first step, a graph of retweet interactions was built from a recorded stream of interactions.
2. In this graph, the largest connected component was determined, as this has the largest impact on the social network.
3. Within the largest connected component, two (polarized) communities were then determined.
4. Within these two communities, characteristic metrics were determined
5. to assign the same community labels at different points in time. This study did this by assigning the colours blue and red. Red represents the community with "unusual" interaction behaviour. Other labels would work as well.



**Figure 3.** Archtitecture used to record German-language Twitter traffic via the Public Streaming API. The Public Streaming API provides a representative sample of about 1% of the total traffic.

This procedure is also explained in Figure 4 using an example graph for illustrative purposes. The reader should consider that millions of interactions were captured in the context of this study, and the graphs become correspondingly complex.



**Figure 4.** Visualization of graph-based determination of echo chambers. An echo chamber is determined in multiple steps. Within the graph, the largest component, within the largest component two polarized communities. The labelling of the communities in red and blue is done by a metric-based property analysis of both communities.

### 2.2.1. Graph Building and Largest Component Detection (see Figure 4 steps ❶ + ❷)

The NetworkX library was used to build retweet interaction graphs. NetworkX [35] is a powerful Python library for analyzing and understanding complex networks in various fields, including social networks, transportation networks, and biological networks. It allows creating of directed and undirected graphs, using various node and edge attributes, and the ability to perform graph algorithms and calculate network statistics. Additionally, NetworkX is built to work with other scientific computing libraries, such as NumPy and SciPy, and can easily integrate with other data analysis and visualization software.

The reader can find the corresponding Python code in Appendix A in Listing 2. Listing 3 shows the code for determining the largest connected component within the graph.

### 2.2.2. Community Detection (see Figure 4 step ❸)

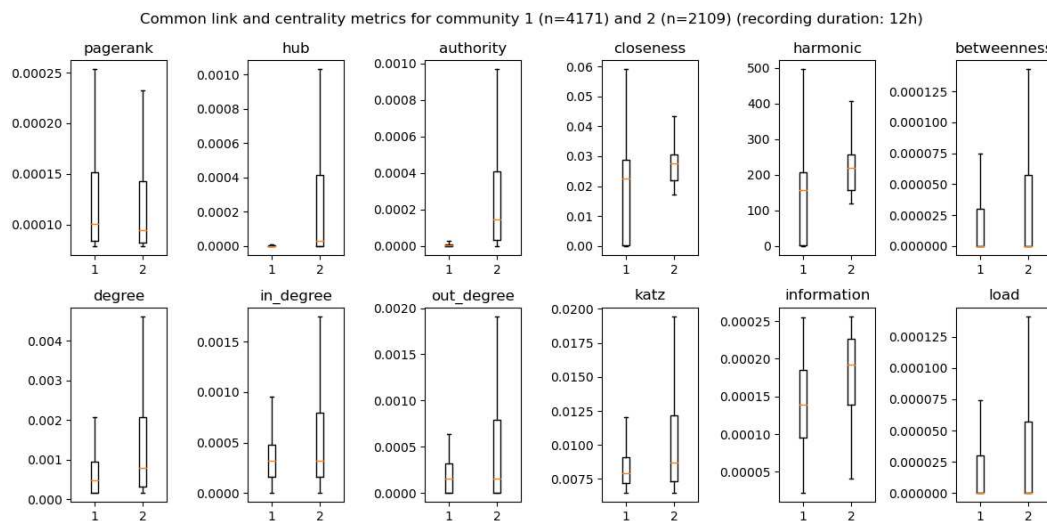
There are many algorithms for community detection in graphs, all of which come with different advantages and disadvantages [24]. It has been found that bipartitions (which can be applied recursively if necessary to identify more than two communities) are particularly suitable for detecting echo chambers in polarized scenarios. In the present study case, detecting two well-separable communities is sufficient. In consequence, algorithms that cannot be parameterized to derive a specific amount of communities in a graph (in our case,  $n = 2$ ) are unsuitable for this use case. Therefore, and according to the literature, primarily the following algorithms can be considered:

- The Kernighan-Lin Bisection algorithm [36] partitions a network into two sets by iteratively swapping pairs of nodes to reduce the edge cut between the two sets.
- The asynchronous fluid communities algorithm [37] is based on the idea of fluids interacting in an environment, expanding and pushing each other. Its initialization is random, so found communities may vary on different executions.
- The Clauset-Newman-Moore greedy modularity maximization algorithm [38] finds a community partition with the largest modularity. Greedy modularity maximization begins with each node in its own community and repeatedly joins the pair of communities that lead to the largest modularity until no further increase in modularity is possible (a maximum). To obtain precisely  $n$  communities, the *cutoff* and *best<sub>n</sub>* parameters can be set to  $n$  (in our case,  $n = 2$ ).

This study used the **greedy modularity maximization** approach. However, other community detection might also work but might show slightly different results. Investigations carried out during this study have not shown any significant differences in the present use case, so it is probably irrelevant which of these algorithms is used. The reader can find the corresponding Python code in Appendix A in Listing 4.

### 2.2.3. Graph Metrics and Colouring (see Figure 4 step ❹)

There are countless metrics to measure graphs [39]. The NetworkX library provides Pagerank, and HITS for link analysis and several centrality measures like closeness-, harmonic-, betweenness-, degree-, Katz-, information-, or load-centrality. This paper will not go into detail on all of these metrics. Still, these metrics were examined to determine which is particularly useful in reliably distinguishing the characteristics of two communities. The reader can find these results in Figure 5.



**Figure 5.** Comparison of different centrality metrics.


Figure 5 compares different centrality measures in the two communities identified using the greedy modularity maximization algorithm. It is striking that the boxes of the boxplots overlap for almost all metrics, which can be problematic for clear assignments. However, the HITS metric (hub + authority) shows almost no deflection for one community and significant deflections for the other. The HITS metrics (especially the authority values) are an excellent discriminator in differentiating the communities. HITS (Hyperlink-Induced Topic Search, [40]) is a link analysis algorithm used to measure the importance of web pages (the Google PageRank algorithm pursues similar goals but shows similar values for both communities and is, therefore, more unsuitable for echo chamber detection). The HITS algorithm was initially developed to sort hits in search engines according to relevance. The algorithm assigns two values to each node, a hub score and an authority score.

- A hub is a node that links to many other nodes. It is a central information point and is used to connect different topics. In this study context, a hub is an account that mainly *retweets* other accounts. So, a hub is a content disseminator.
- An authority is a node that is linked to many other nodes. It can be seen as an authoritative source of a community used to get information on a particular topic. In this study, an authority is an account mainly *retweeted* by other accounts. So, an authority is a content provider.

The community with higher HITS authority values is referred to as the red community, and the other as the blue community. So, both communities clearly show different behaviour in how the content of accounts is retweeted. This is much more pronounced in the red community than in the blue community. This becomes even clearer when the retweet graph is plotted using an Atlas Force algorithm [28], and the node size is set proportional to the HITS authority values. Figure 6 clearly shows that the blue community looks "normal" and the red community has exceptionally many "authorities". These "authorities" fuel the red echo chamber.

The reader can find the corresponding Python code in Appendix A. Listing 5 shows the metrics calculation, and Listing 6 shows how the communities are labelled in the graph using colour labels.

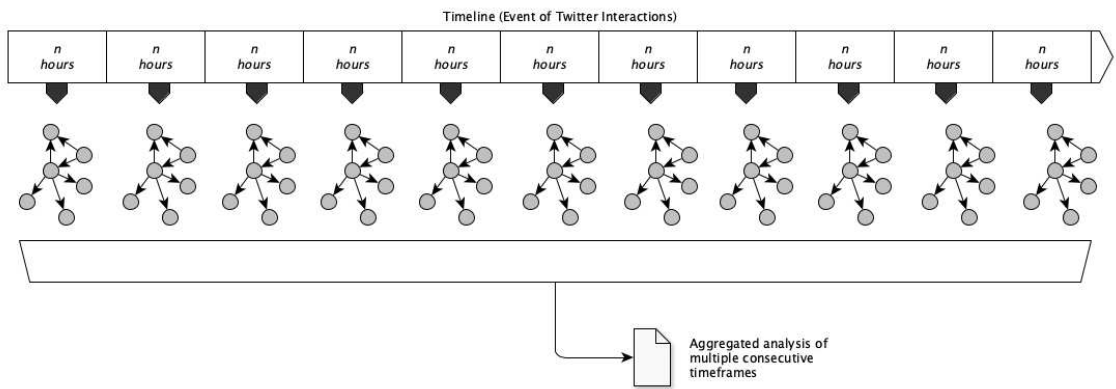




**Figure 6.** The visualization of an echo chamber in a graph recorded over 6 hours. The node size of each observed account is set according to the authority metric of the observed accounts. It can be seen that accounts with high authority values are predominantly located in the red community (the echo chamber).

2.3. Evaluation

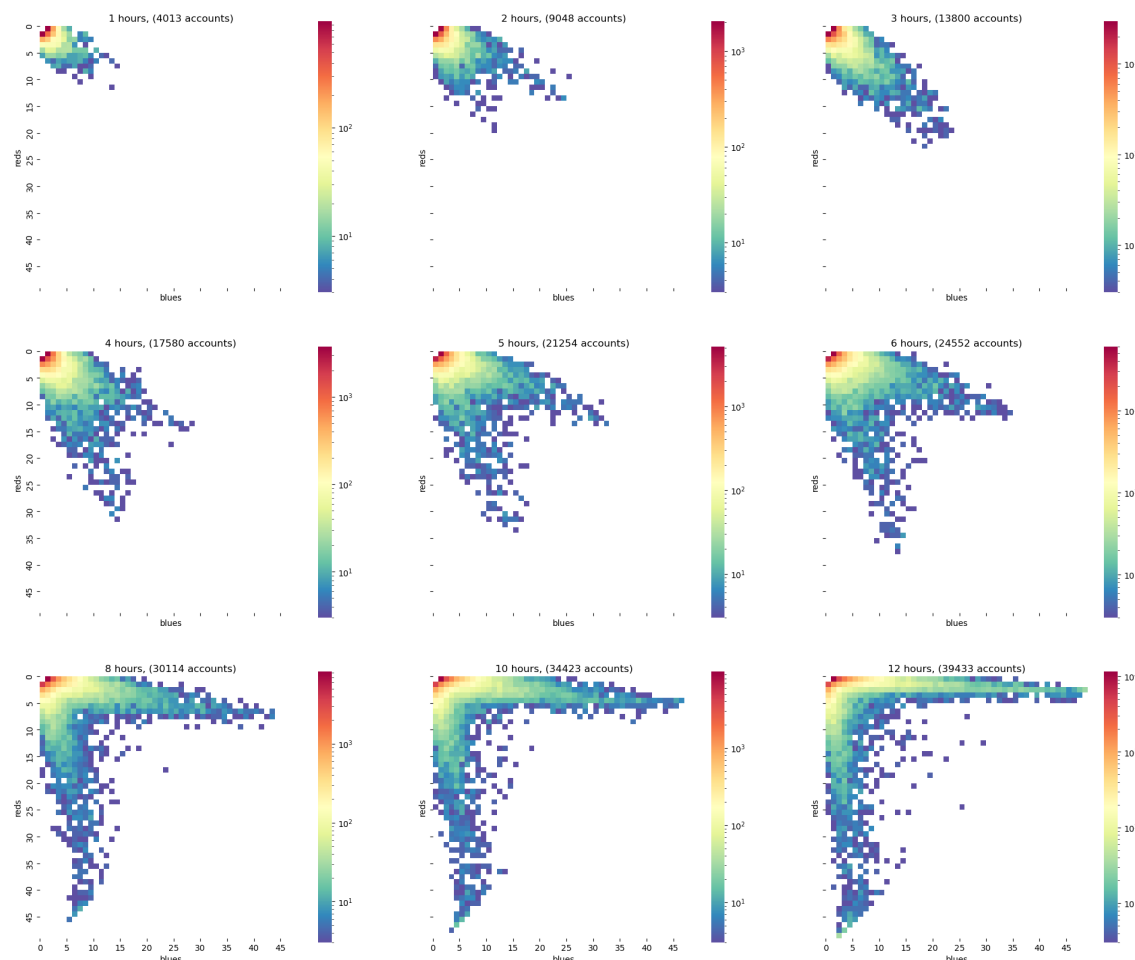
As part of the evaluation, the extent to which the proposed method for detecting echo chambers would work in practice had to be investigated. As there is no common definition of an echo chamber in the literature, there are hardly any reference data sets with which the presented method can be compared. Therefore, individual accounts’ red/blue allocation was examined over time. The evaluation, therefore, focuses more on the aspect of how stable the red/blue classifications of an account are over time. This was done as shown in Figure 7.



**Figure 7.** Principle of how the Echo Chamber evaluation was conducted. Consecutive timespans were taken in each case. For each of these periods, a graph was formed and evaluated. The totality of the results of all graphs was evaluated within the framework of the evaluation also to be able to collect statistical fluctuation ranges.

A retweet graph was created for a period of  $n$  hours, and the red and blue communities were determined as shown in Sections 2.2.1, 2.2.2 and 2.2.3. This graph was compared to the graph of further  $m$  previous consecutive periods of  $n$  hours to determine how stable the red and blue assignments were for individual accounts. The evaluation examined 45 time periods between 1 and 24 hours. Of

course, it is unlikely that individual accounts can be unambiguously assigned to only one of the two communities (some accounts can share the positions of both communities). For stable echo chamber detection, it is necessary to recognize clear red and blue domains, which statistically indicate that predominantly red or blue content is shared. The result of this analysis can be found in Figure 8, which shows that the duration of a period plays a significant role. The distributions for one hour to three hours (first row of Figure 8) have a more circular or oval character and are formed along the main diagonals. From 4 to 6 hours of recording time (second row of Figure 8), the reader observes how two fingers are formed to the left and right of the main diagonal, which correspond to a red and a blue community.



**Figure 8.** Visualization of 45 consecutive time frames of different length (1 to 12 hours). It can be clearly seen that a sufficient recording time is required to achieve clear classifications of individual accounts. The longer a recording time frame is, the more clearly accounts can be assigned to red and blue communities (poorly classifiable accounts are shown in the middle of the diagrams).

For 8 to 12 hours, these fingers start to nestle against the axes (third row of Figure 8). The central area is hardly populated. Thus, observing a sufficiently long interval can record enough interactions to identify a red and a blue community reliably. However, the time interval must not be too short. This period also depends on how many interactions are generated in the first place (in Spanish and English, for example, one will be able to capture much more retweets in 12h than in German, simply because there are more English- and Spanish-language tweets).

In particular, since the classification into red and blue areas is based solely on interactions and not content, little can be said about the nature of the red and blue communities, except that the red community shows strikingly different retweet behaviour than the blue community. Why this is so,

and what content is distributed in such a strikingly different way, was based on a preliminary content analysis of the accounts that achieved the highest authority metric scores overall observed periods. These accounts are retweeted most frequently and - based on the observed interaction patterns - are considered "trusted sources" or opinion leaders for the respective communities and thus set the "tone in the community".

So, for this study was decided to work with a time span of 24 hours and 60 consecutive time frames which leads to a stable classification for plenty of accounts (see Figure 9 in Sec. 3). The results are shown in the following Sec. 3.


### 3. Results

The presented results are based on an analysis of 60 consecutive 24-hour periods. Thus, the evaluation covers about two months of the public stream sample of the German-speaking Twitter space. About 180,000 accounts and their observed retweet interaction patterns are considered. These accounts became active in the blue or red community and retweeted red or blue content. Strictly this database is used to answer the two research questions formulated.

#### 3.1. Research Question 1: Classification results

The primary question of this study was whether it is possible to perform echo chamber detection using graph-based analysis based on a sample of network interactions without the need for content analysis. This could be answered positively, but one must have sufficient interactions for evaluation. In the case of the German language area, it has been shown that one must record at least 12 hours of Twitter retweet interactions for this (see Figure 7). More is better, but also more computationally intensive.

Figure 9 shows the evaluation for the German-speaking area for 60 consecutive 24-hour periods, i.e. about two months. One can see that the majority of accounts appear only a few times. This is a common phenomenon in scale-free social networks. Nevertheless, there is an accumulation along the red and blue axes. The distribution of retweets is, therefore, not equally distributed between the two communities. If this were the case, there would be an accumulation along the main diagonal.



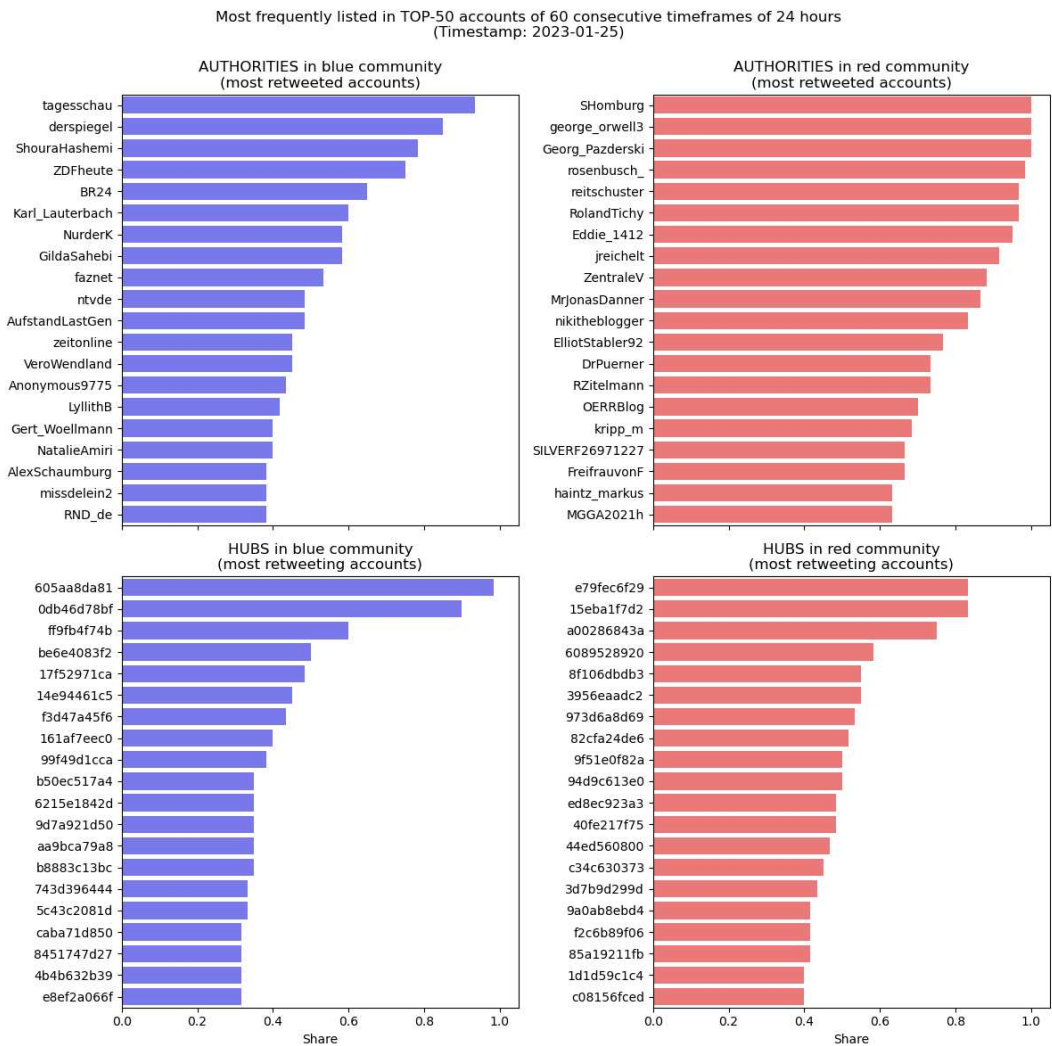
**Figure 9.** Visualization of the frequency of blue/red classifications of about 180000 Twitter accounts in two months. A 24h time window was always used to classify an account as red or blue.

Therefore, the proposed approach is suitable for echo chamber detection. Nevertheless, the classification does not tell us much more than that the red community differs from the blue community by a significantly different distribution of authority indicators (HITS metric). Since no content-related evaluation of the shared content was carried out, little can be said about the content of both communities. The content could include anything, e.g. Mickey Mouse films, make-up tips, sports news, political content, health tips or completely other content. It was impossible to carry out a comprehensive and detailed analysis of all 180,000 Twitter accounts covered. Nevertheless, an assessment of the content of the red and blue communities is of course, of interest. Therefore, we determined the accounts with the highest authority values in each of the 60 24-hour periods and counted how often they appeared in the TOP-50 accounts. The higher an account appears, the more regularly it is retweeted and the more content-related influence it has on the respective community.


Figure 10 shows the result of this analysis for the blue and red communities. Tables A1 and A2 in Appendix B list further information per account, such as the description, the verification status (Twitter blue tick), or content categories. A verified account on Twitter is an account with a blue badge next to the account name, indicating that Twitter has confirmed that the account is of public interest and is the authentic account of the person or entity it represents. This verification helps users differentiate between authentic accounts and impostor accounts. To become verified on Twitter, users must provide information about their identity and account history to demonstrate that they meet Twitter's criteria for verification.

Readers who regularly follow the media in the German-speaking world will recognise at first glance that the red accounts are strongly right-wing populist or critical of Corona measures. In contrast, the blue accounts are more influenced by classic TV or print media. This gives a first impression of the way of communication in both communities. A closer look at the red accounts further reveals that some of these accounts spread notable Russian narratives up to the cheapest Russian propaganda in the context of the Russian war of aggression on Ukraine launched on 24 February 2022.

Figure 11 shows quite different types of accounts in the highest-ranked authority accounts. While only about 25% verified accounts characterize the red community, the blue community has 55%. A similar picture emerges for journalistic accounts. Here, too, only about 20% of the accounts in the red community have a journalistic background, half as many as in the blue community. Accounts of politicians, on the other hand, are equally strong (or weak) in both communities, at only 5%. On the other hand, Troll-like accounts do not occur in the blue community but form the majority of opinion-shaping accounts in the red community. Such Troll-like accounts are false online identities often controlled by Russian government agencies or their supported organizations [41]. They are used to influence online discussions and spread certain political narratives. Since the annexation of Crimea in 2014, these troll accounts have been increasingly used as part of the Russian information war to influence and destabilize public opinion in Western countries. It can therefore be stated that verified and often journalistic accounts characterize the blue community. On the other hand, the red community is characterized by accounts with a solid troll-like character (even if this cannot be proven beyond doubt in the context of this study).



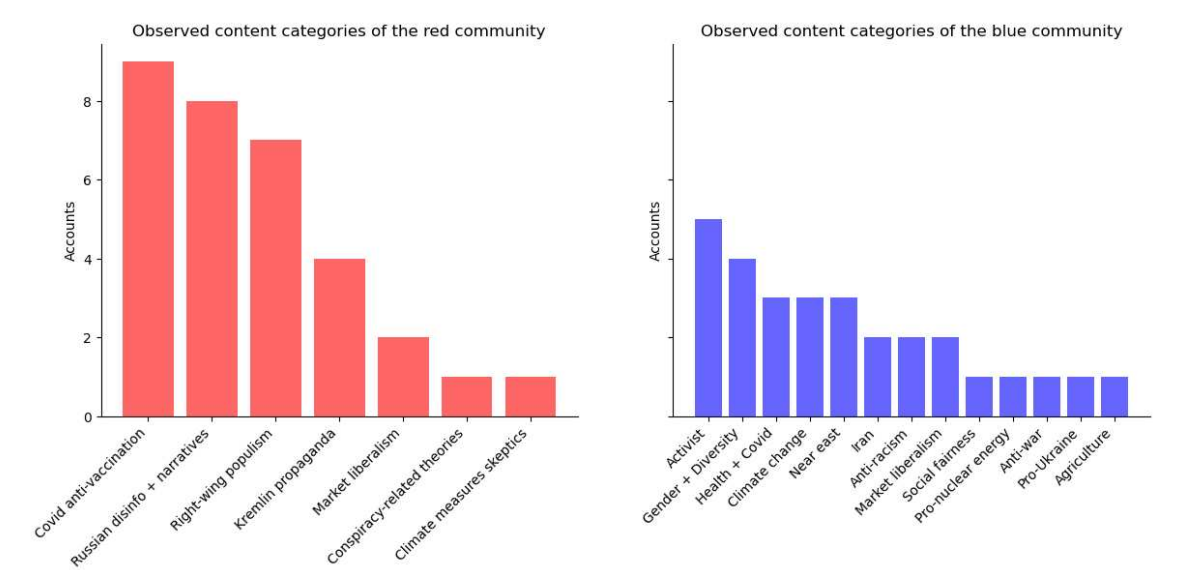
**Figure 10.** Most frequently observed accounts in the TOP-50 hub and authority accounts of the red and blue communities in 60 consecutive time frames. Hub accounts are displayed anonymously to protect personal rights. Due to the accounts’ range, frequency and professionalism, authority accounts must have a professional awareness of broadcasting, so anonymization is dispensed with. Important: The accounts’ categorizations were derived from posting behaviour, not based on content analyses. The red/blue categorization may therefore be debatable in individual cases.



**Figure 11.** Content analysis of the accounts with the highest Authority values in the blue and red communities. These accounts are frequently retweeted and therefore have a formative influence on the community. Accounts can be assigned to multiple categories (e.g. verified journalistic accounts). Other uncategorized accounts may exist if the totals do not add up to 100%.



Figure 12 shows the content analysis for the months December 2022 and January 2023. It can be seen that the content of the red community is focused on fewer topics. The content dissemination is dominated by accounts that oppose Covid vaccination, often share right-wing populist content, and share Russian narratives (e.g., Russia was forced into war) and disinformation (e.g., Ukraine oppresses its Russian-speaking population). This triad of anti-Covid, right-wing populism, and pro-Russian positions seems particularly pronounced in the red community. In contrast, the blue community is thematically much more heterogeneous, although several activist accounts also characterize it. Here, the topics range from Iran, gender + diversity, Covid, climate change, anti-racism, and social fairness to pro-nuclear energy and agriculture. Only in the case of market liberalism is a slight overlap between the two communities.



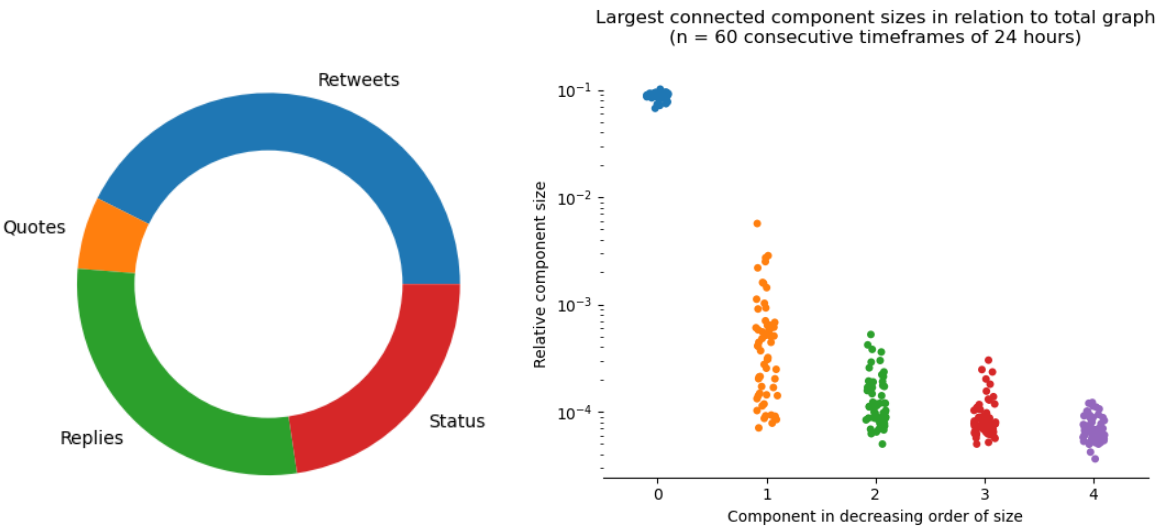
**Figure 12.** Content analysis of the accounts with the highest Authority values in the blue and red communities. These accounts are frequently retweeted and therefore have a formative influence on the community.

3.2. Research Question 2: Scope and quantity considerations

The second research question focuses on how much data is taken into account and for which scopes of the interaction graph can be reliably made conclusions. Figure 13 shows the observed ratio of retweets, quotes, replies and status posts on the left. Retweets account for just over a third of all interactions. This third of the traffic has the advantage that it can be processed without computational intensive content analysis, i.e. complex NLP procedures.


The second question is how large the proportion of the largest connected component is within the overall retweet graph. Figure 13 (right side) shows this ratio for all 60 24h-graphs of the observation period. According to this, the largest connected component makes up about 10% of the full retweet graph. The second largest connected component only accounts for a maximum of 1%, and so on<sup>1</sup>.

<sup>1</sup> The reader should consider that only accounts retweeted at least once were counted as part of the largest connected component. If we also include accounts that have not been retweeted once, the largest connected component would make up about 80% of the total graph (not shown in Figure 13). However, most of these accounts would be "dead and inactive" endpoints.



**Figure 13.** Visualization of connected component sizes. On the left side, one can see the observed share of Twitter interactions (retweet, quote, reply, status post). Only retweet interactions have been used in this study. On the right side, one can see the size differences of the largest (n=0), second largest (n=1), and so on the connected components of the retweet graph.

The reader should be aware that the presented approach derives insights on only about 1/3 of the available data (retweets only), enabling categorizations of about 10% of the observed retweet network. Therefore, few conclusions about the bulk of the network’s interactions and accounts can be drawn. Estimates refer to only about 10% of the total network, of which about 1/3 show "unusual" communication patterns. A typical ratio between the red and blue portions in the largest connected component is shown in Figure 14.



**Figure 14.** Exemplary evaluation of a Largest Connected Component of the serial examination. Typically, the ratio of red to blue accounts is about 1 to 3, so red nodes make up only about 1/3 of the accounts, but share a disproportionate amount of content and therefore dominate the attention of many content recommender algorithms.

4. Discussion

So, the reader sees that there are also drawbacks to the presented structural analysis approach. One limitation is that structural analysis alone may not provide insight into the content or dynamics of echo chambers, so other methods, such as natural language processing, automatic sentiment, or

even manual content analysis, may also be needed. The choice of community detection algorithm or centrality measure can also affect the results obtained.

#### 4.1. Limitations

Graph-based analysis methods enabled the red and blue groups to be distinguished from each other in this study. However, which content is primarily exchanged in these echo chambers remains fundamentally unclear. In this study, we did not examine this content in detail, only for the most significant accounts with the highest authority values in the red and blue groups. Yet, it appears that the blue group makes up the majority with the usual political-content differences. Based on the authorities, the red group is more likely to share Corona denying, alt-right theory, xenophobia and strongly right-wing populism content. In addition, pro-Russian narratives seem to be spread more and more, and some accounts give the impression of being run by Russian troll factories. However, the reader should be aware that this content analysis is somewhat preliminary and has not been carried out according to recognized scientific standards. But the differences in content between the two groups seem astonishingly clear. Furthermore, only the German-language Twitter datastream was examined. Further studies should also study whether the presented approach can be applied to other language areas.

Most likely, the author assumes, the method loses its power when applied to several language areas simultaneously because connected components are then likely to form along the language areas. The polarity approach should thus reach its limits. Large and international language spaces such as English or Spanish could also become problematic, as these languages are used in different countries, in different time zones and very different social contexts (e.g. in Puerto Rico, Argentina and Spain). English, the leading world language, is used internationally and is no longer oriented towards clearly definable language spaces.

Although, the approach can derive a blue and a red group. Both groups are deduced exclusively based on a HITS metric. The fact that the red group mainly contains content that can be assigned to the alt-right spectrum may be pure coincidence or may only apply to the German language area. It could also be that these communities are systematically undermined by Russian troll activities and do not provide an accurate picture. This is up for further investigation.

#### 4.2. Threats to validity

The results of this study may be affected by internal and external threats to validity. Validity is an essential concept in data science studies as it determines the accuracy of the results obtained from the study. Validity refers to the extent to which a study measures what it is intended to measure. Internal and external threats to validity can affect the accuracy of a study's results. Internal threats arise within the study, such as selection or instrumentation bias. External threats, such as media topics and political or other events, arise outside the study and are not entirely under the experimenter's control. The following sections, therefore, address how the aspects and effects of both threats to validity were considered and minimized.

##### 4.2.1. Internal Threats

**Selection Bias:** Selection bias occurs when participants are not randomly selected for a study, resulting in a sample that does not accurately represent the studied population. This study deliberately worked with a sample of the complete German-speaking area. Studies have shown that the public Twitter streaming API provides representative data [23], so it can be assumed that this study has largely eliminated this bias.

**Instrumentation Bias:** Instrumentation bias occurs when instruments used to measure variables are not reliable or valid. This study successfully applied a recording solution that has been used in

other peer-reviewed studies [31]. The solution is used to capture a monthly updated sample of the German Twitter stream [42]. A variety of researchers use this dataset<sup>2</sup>.

**History Bias:** History bias occurs when events occurring before or during a study influence its results. The reader should be aware that the content analysis of Section 3.1, in particular, should be viewed against the background of events in December 2022 and January 2023. The study makes no claims that the procedure can identify right-wing populist themes, for example. The fact that these themes appear in the analysis has to do with the historical circumstances of the analysis period.

**Maturation Bias:** Maturation bias occurs when changes in participants' behaviour over time influence results obtained from a longitudinal study. This bias should not have occurred because the study's findings were not used to influence any behaviour. In other words, the analyzed accounts did not know that they were part of the study and, therefore, could not adjust their behaviour.

#### 4.2.2. External Threats

**Construct Validity:** Construct validity refers to whether a measure accurately reflects what it is intended to measure (i.e., its construct). Within the framework of this study, it was possible to identify accounts in the right-wing populist community. Accounts that have a strong troll-like character and are presumably Kremlin-orchestrated also appear. This is undoubtedly an effect of the historical context (Russia-Ukraine war) and the general political mood in Germany, which is characterised by increasing right-wing populist divisive movements. The study does not state that these effects exist to the same extent in other language areas. Nor is it possible to make any statement about how the process would behave if these effects no longer existed.

**External Validity:** External validity refers to whether findings from one sample can be generalized across other samples. This study examined the results in the context of a study lasting several months in German-speaking countries. Whether the results are transferable to shorter or longer periods and other language areas must be investigated in further studies.

## 5. Conclusions

This study investigated the use of purely structural analysis of communication and interaction patterns to identify echo chambers in a sample of the German-speaking Twitter network. The results suggest that the greedy modularity maximization and HITS metric are the most effective methods for identifying echo chambers. The analysis covers 60 consecutive 24-hour periods and considers 180,000 accounts and their observed retweet interaction patterns. The results show that a red community (the echo chamber) is focused on fewer topics with a triad of anti-Covid, right-wing populism, and pro-Russian positions. In contrast, the blue community is more heterogeneous. However, it should be noted that the approach provides insights on only about 1/3 of the available data (retweets only) and categorizations of about 10% of the observed retweet network, limiting the conclusions that can be drawn about the bulk of the network's interactions and accounts. So few conclusions can be made about most of the network's interactions and accounts.

Nevertheless, these analyses help to identify "problematic interaction patterns" in social networks with little effort. Social networks could, for example, give such echo chamber accounts less weight in their content recommender algorithms, as these accounts often only repeat the same content again and again. This could also minimise the misuse of social networks to spread disinformation by problematic actors (such as Russian disinformation bot networks). It is essential to understand that not the content but only the interaction patterns would be used as a decision criterion. This is, therefore, not censorship of content, which would also be problematic from the author's point of view.

---

<sup>2</sup> The author approved more than 100 access requests to this dataset so far.

**Funding:** This research received no external funding.

**Institutional Review Board Statement:** Not applicable.

**Data Availability Statement:** The data used in this study is provided as a continuously updated Zenodo repository [42]. Due to data protection requirements and Twitter usage agreements, the data can only be requested and released for members of scientific institutions and for scientific research questions. Restrictions on the use of the data must be considered by data users.

**Acknowledgments:** The author would like to express his gratitude to the reviewers for their valuable feedback and time in reviewing the manuscript. Their insights have helped improve the quality. The author is grateful for their contributions to the scientific community.

**Conflicts of Interest:** The author declares no conflict of interest.

## Appendix A

### Listing 1: Recording code

---

```
import tweepy
import os
import time
from pymongo import MongoClient
from dateutil import parser

TRACKS      = os.environ.get('TRACK', "a,b,c,...,x,y,z").split(",")
LANGUAGES    = os.environ.get('LANGUAGES', 'de').split(",")
MONGO_SVC    = os.environ.get('MONGO_SVC', 'host.docker.internal')
MONGO_SVC_PORT = int(os.environ.get('MONGO_SVC_PORT', '27017'))
RECOVERY     = int(os.environ.get('RECOVERY', '30'))
COLLECTION   = os.environ.get('COLLECTION', 'tweets')

client = MongoClient(MONGO_SVC, MONGO_SVC_PORT)
db = client['twista']
collection = db[COLLECTION]

try:
    collection.create_index("created_at")
except Exception as ex:
    print(f"Index already exists {ex}")

creds = {
    "API_KEY": os.environ.get('API_KEY'),
    "API_SECRET_KEY": os.environ.get('API_SECRET_KEY'),
    "ACCESS_TOKEN": os.environ.get('ACCESS_TOKEN'),
    "ACCESS_TOKEN_SECRET": os.environ.get('ACCESS_TOKEN_SECRET'),
}

class Recorder(tweepy.Stream):
    def on_tweet(self, status):
        self.on_status(status)

    def on_status(self, status):
        try:
            tweet = status._json # status.data
            tweet['created_at'] = parser.parse(tweet['created_at'])
            collection.insert_one(tweet)
            print(status.text)
        except Exception as ex:
            print(f"Exception {ex}")

recorder = Recorder(
    creds['API_KEY'],
    creds['API_SECRET_KEY'],
    creds['ACCESS_TOKEN'],
    creds['ACCESS_TOKEN_SECRET']
)
```



```

)

while True:
    try:
        recorder . filter (track=TRACKS, languages=LANGUAGES)
    except Exception as ex:
        print (f"Argh, getting this exception : {ex}")
        print ("Disconnecting stream")
        recorder . disconnect ()
        print (f"Waiting some {RECOVERY} seconds to respect Twitter (re-)connection rate limits")
        time . sleep (RECOVERY)

```

---

### Listing 2: Graph building code

---

```

# Retweet Interaction Graph Building
def graph(db : Collection , start : datetime, end : datetime) -> nx.DiGraph:
    G = nx.DiGraph()

    query      = { "retweeted_status": { "$ne": None }, "created_at": { "$gte": start , "$lt": end }}
    projection = { "user.screen_name": 1, "retweeted_status.user.screen_name": 1 }

    for doc in db.find(query, projection):
        o = doc["retweeted_status"]["user"]["screen_name"]
        r = doc["user"]["screen_name"]
        G.add_node(o)
        G.add_node(r)
        G.add_edge(r, o)

    return G

```

---

### Listing 3: Largest connected component detection

---

```

# Largest Connected Component Detection
def lcc(G : nx.DiGraph, predecessors=1, sizes=False) -> nx.DiGraph:
    nodes = [n for n in G.nodes() if len(list(G.predecessors(n))) >= predecessors]
    F = G.subgraph(nodes)
    graphs = [F.subgraph(c) for c in sorted(nx.connected_components(F.to_undirected()), key=len, reverse=True)]
    if not graphs:
        return nx.DiGraph()
    if sizes:
        return [len(g) / len(G) for g in graphs]
    return graphs[0]

```

---

### Listing 4: Community detection

---

```

# Community Detection
def community(g : nx.DiGraph, n=2) -> nx.DiGraph:
    for i, c in enumerate(nx.algorithms.community.greedy_modularity_communities(g, best_n=n)):
        sg = g.subgraph(c)
        nx.set_node_attributes(sg, i, name="community")
    return g

```

---

### Listing 5: Metrics calculation

---

```

# Metrics calculation
def metrics(g : nx.DiGraph) -> nx.DiGraph:
    h, a = nx.hits(g)
    nx.set_node_attributes(g, h, name="hub")
    nx.set_node_attributes(g, a, name="authority")
    # Optional, just used for metrics comparison
    #nx.set_node_attributes(g, nx.degree_centrality(g), name="degree")
    #nx.set_node_attributes(g, nx.in_degree_centrality(g), name="in_degree")

```

```
#nx.set_node_attributes(g, nx.out_degree centrality(g), name="out_degree")
#nx.set_node_attributes(g, nx.closeness centrality (g), name="closeness")
#nx.set_node_attributes(g, nx.harmonic centrality(g), name="harmonic")
#nx.set_node_attributes(g, nx.pagerank(g), name="pagerank")
#nx.set_node_attributes(g, nx.information centrality (g.to_undirected(as_view=True)), name="information")
#nx.set_node_attributes(g, nx.betweenness centrality(g), name="betweenness")
#nx.set_node_attributes(g, nx.katz centrality (g), name="katz")
#nx.set_node_attributes(g, nx.load centrality (g), name="load")
return g
```

Listing 6: Colouring

```
# Colouring
def colorize (g : nx.DiGraph, percentile=75) -> : nx.DiGraph:
    a = [(n, d) for n, d in g.nodes().items() if d['community'] == 0]
    b = [(n, d) for n, d in g.nodes().items() if d['community'] == 1]

    na = [n for n, _ in a]
    nb = [n for n, _ in b]

    m_a = np.percentile([d['authority'] for _, d in a], percentile)
    m_b = np.percentile([d['authority'] for _, d in b], percentile)

    nx.set_node_attributes(g.subgraph(na if m_a > m_b else nb), "red", name="color")
    nx.set_node_attributes(g.subgraph(nb if m_a > m_b else na), "blue", name="color")

    return g
```

Appendix B

**Table A1.** List of 20 accounts with the highest Authority scores (HITS metric) recorded in the TOP-50 accounts of the red community in 60 consecutive time frames, sorted in descending order by frequency. These accounts are retweeted particularly frequently.

Account	Verified	Description	Color	Topics
SHomburg	No	Professor, @UniHannover, Autor "Corona-Getwitter", Leseprobe: <a href="http://bit.ly/3H9Jlil">http://bit.ly/3H9Jlil</a> Bildung statt Haltung, Like ≠ Zustimmung	red	Covid anti-vaccination
george_orwell3	No	Mit fast tödlicher Sicherheit bewegen wir uns auf ein Zeitalter totalitärer Diktaturen zu.	red	Russian Misinfo, Kremlin propaganda, Covid anti-vaccination
Georg_Pazderski	Yes	Oberst i.G. (a.D.), Ehemaliger AfD-Fraktionsvorsitzender im Abgeordnetenhaus von Berlin (Follow und RT sind keine Zustimmung)	red	Politician, Right-wing populism
rosenbusch_	No	Independent Journalist/photographer, committed to the truth to the best of my knowledge and belief since 2001.	red	Russian Misinfo, Covid anti-vaccination

Table A1. Cont.

Account	Verified	Description	Color	Topics
reitschuster	Yes	In 16 Jahren als Korrespondent in Moskau allergisch geworden gegen Ideologen, Sozialismus-Nostalgiker und Journalisten-Kollegen, die die Regierung loben.	red	Covid anti-vaccination
RolandTichy	Yes	Roland Tichy, Gründer TichysEinblick; Das Leben ist kein Ponyhof   Impressum: <a href="https://tichyseinblick.de/impressum/">https://tichyseinblick.de/impressum/</a>	red	Journalist, Market Liberalism
Eddie_1412	No	Was Lucky Luke im Wilden Westen bin ich auf Twitter. Ich blockiere schneller als mein Schatten...	red	Kremlin propaganda, Right-wing populism, Russian Misinfo
jreichelt	Yes	I'm back!	red	Online-Journalist, Right-wing populism
ZentraleV	No	Zentrale Ermittlungsstelle fremde Medien Ost	red	Russian Disinfo, Kremlin propaganda
MrJonasDanner	No	Der Account, vor dem euch die Tagesschau immer gewarnt hat.	red	Russian Misinfo, Covid anti-vaccination
nikitheblogger	No	YouTuber mit 300.000 Abonnenten, freier Journalist und politischer Blogger	red	Online-Journalist, Right-wing populism
ElliotStabler92	No	Marcel	red	Conspiracy-related theories, Russian Misinfo
DrPuerner	No	Facharzt für Öffentliches Gesundheitswesen, Epidemiologe, Impfstatus: privat	red	Medical, Covid anti-vaccination
RZitelmann	No	ILOVECAPITALISM Kapitalismus ist nicht das Problem, sondern die Lösung. Official German profile	red	Market Liberalism, Climate measures skeptics
OERRBlog	Yes	Kritische Beobachter des deutschen Öffentlich-rechtlichen Rundfunks. Für eine Verkleinerung und Kostenreduzierung. ReformOerr	red	Media blog
kripp_m	No	Eine Investition in Wissen bringt immer noch die besten Zinsen. (Benjamin Franklin)	red	Russian Misinfo, Covid anti-vaccination
SILVERF26971227	No	VOX POPULI - VOX DEI	red	Russian Disinfo, Kremlin propaganda
FreifrauonF	No	Natürlich Patriotin! Für ein Europa der Vaterländer! Du kannst alles verkaufen, nur nicht deine Leute und nicht dein Land. Pronomen: Mom/Milf	red	Right-wing populism
haintz_markus	No	Rechtsanwalt und Journalist FreeAssange	red	Right-wing populism, Russian Misinfo, Covid anti-vaccination
MGGA2021h	No	Für ein starkes Europa mit eigenen Nationen. Für glückliche+ungeimpfte Kinder. Keine Islamisierung. "Journalist". Master o.t.U.	red	Right-wing populism, Russian Misinfo, Covid anti-vaccination

**Table A2.** List of 20 accounts with the highest Authority scores (HITS metric) recorded in the TOP-50 accounts of the blue community in 60 consecutive time frames, sorted in descending order by frequency. These accounts are retweeted particularly frequently.

Account	Verified	Description	Color	Topics
tagesschau	Yes	Schlagzeilen von <a href="https://tagesschau.de">https://tagesschau.de</a>	blue	TV news, public service
derspiegel	Yes	Nachrichten, Analysen, Kommentare, Videos, Podcasts: Mehr als 500 SPIEGEL-Journalistinnen und Journalisten decken auf, setzen Themen und sorgen für Kontext.	blue	Print media
ShouraHashemi	No	Bitte googelt mich nicht. she/her	blue	Iran, Activist
ZDFheute	No	Hier twittert die ZDFheute-Redaktion Nachrichten, Videos und Hintergründe.	blue	TV news, public service
BR24	Yes	Hier ist Bayern.	blue	TV station, public service
Karl_Lauterbach	Yes	SPD Bundestagsabgeordneter, Bundesgesundheitsminister, der hier selbst und privat tweetet.	blue	Politician, Health, Covid
NurderK	No	Skills can be taught. Character you either have or you don't have.	blue	Climate change, Social fairness
GildaSahebi	Yes	Journalistin / Ärztin / Politikwissenschaftlerin, @tazgezwitscher, Antisemitismus, Rassismus, Naher Osten, Wissenschaft, Medizin.	blue	Activist, Health, Anti-Racism, Near east
faznet	Yes	Die wichtigsten Nachrichten des Tages, die besten Faz.net-Artikel und Empfehlungen der Redaktion.	blue	Print media
ntvde	Yes	NTV Nachrichten: FürAlledieFragenhaben, Mehr von ntv: @ntv_EIL, @ntvde_politik, @teleboerse, @ntvde_Sport, @ntvde_Auto, @ntvpodcast	blue	TV news, private broadcaster
AufstandLastGen	No	Wir sind die LetzteGeneration, die den völligen Klimakollaps noch aufhalten kann!	blue	Activist, Climate Change
zeitonline	Yes	Ja, das ist unser offizieller Twitter-Account. Hier bekommen Sie die wichtigsten Geschichten und aktuelle News.	blue	Print media
VeroWendland	No	Energoblogger. Ecomodernist. Science, Technology, Society Studies. Eastern Europe. Status: Reaktorversteherin.	blue	Climate change, Energy transition, Pro-nuclear energy
Anonymous9775	No	Hier twittert Anonymous gegen Faschismus, Rassismus, Ungerechtigkeit, Unterdrückung, Zensur, Kriege, Diktatoren, Sekten	blue	Activist, Anti-Racism, Anti-War
LyllithB	No	nix mit Medien, Orwell, Psychedelic, Coffee, Tea, Seventies, HippieGoth, Sea, Regen, Pfalz, Yellow, Cohen, Dylan, Reed, Wader	blue	Influencer, Diversity, Gender
Gert_Woellmann	Yes	Landesvorstand der @fdphh, Kreisvorstand FDP Alstertal-Walddörfer	blue	Politician, Liberalism, Pro-Ukraine
NatalieAmiri	Yes	Int. Correspondent - 2015-20 Iran/Tehran ARD, @DasErste, German Television bureau chief - Anchorwoman @Weltspiegel_ARD, Autorin: "Zwischen den Welten"	blue	Journalist, Iran, Near East
AxelSchaumburg	No	Deziiert liberal, Gegen Rechts- und Linksextremismus und alles Totalitäre, My only "phobias" are Naziphobia und Wokophobia, Following ≠ endorsement	blue	Entrepreneur, Agriculture, Forestry, Liberalism
missdelein2	No	Madeleine	blue	Activist, Gender, Diversity
RND_de	Yes	Exklusive News, interessante Hintergründe und fundierte Meinungen. Hier twittert das RedaktionsNetzwerk Deutschland (RND).	blue	Print media

## References

1. KhosraviNik, M. Right wing populism in the west: Social media discourse and echo chambers. *Insight Turkey* **2017**, *19*, 53–68.
2. Boulianne, S.; Koc-Michalska, K.; Bimber, B. Right-wing populism, social media and echo chambers in Western democracies. *New media & society* **2020**, *22*, 683–699.
3. Margetts, H. Rethinking democracy with social media. *Political Quarterly* **2018**, *90*.
4. Persily, N.; Tucker, J.A. *Social media and democracy: The state of the field, prospects for reform*; Cambridge University Press, 2020.
5. Barberá, P. Social media, echo chambers, and political polarization. *Social media and democracy: The state of the field, prospects for reform* **2020**, *34*.
6. Ross Arguedas, A.; Robertson, C.; Fletcher, R.; Nielsen, R. Echo chambers, filter bubbles, and polarisation: A literature review, 2022.
7. Grömping, M. 'Echo chambers' partisan facebook groups during the 2014 Thai election. *Asia Pacific Media Educator* **2014**, *24*, 39–59.
8. Hayat, T.; Samuel-Azran, T. "You too, second screeners?" Second screeners' echo chambers during the 2016 US elections primaries. *Journal of Broadcasting & Electronic Media* **2017**, *61*, 291–308.
9. Guo, L.; A. Rohde, J.; Wu, H.D. Who is responsible for Twitter's echo chamber problem? Evidence from 2016 US election networks. *Information, Communication & Society* **2020**, *23*, 234–251.
10. Hänska, M.; Bauchowitz, S., Tweeting for Brexit: how social media influenced the referendum. In *Brexit, Trump and the Media*; abramis academic publishing: Bury St Edmunds, UK, 2017; pp. 31–35.
11. Bastos, M.; Mercea, D.; Baronchelli, A. The geographic embedding of online echo chambers: Evidence from the Brexit campaign. *PloS one* **2018**, *13*, e0206841.
12. Gorodnichenko, Y.; Pham, T.; Talavera, O. Social media, sentiment and public opinions: Evidence from# Brexit and# USElection. *European Economic Review* **2021**, *136*, 103772.
13. Badawy, A.; Ferrara, E.; Lerman, K. Analyzing the digital traces of political manipulation: The 2016 Russian interference Twitter campaign. 2018 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM). IEEE, 2018, pp. 258–265.
14. Zannettou, S.; Caulfield, T.; De Cristofaro, E.; Sirivianos, M.; Stringhini, G.; Blackburn, J. Disinformation warfare: Understanding state-sponsored trolls on Twitter and their influence on the web. Companion proceedings of the 2019 world wide web conference, 2019, pp. 218–226.
15. Golovchenko, Y.; Buntain, C.; Eady, G.; Brown, M.A.; Tucker, J.A. Cross-platform state propaganda: Russian trolls on Twitter and YouTube during the 2016 US presidential election. *The International Journal of Press/Politics* **2020**, *25*, 357–389.
16. Karami, A.; Lundy, M.; Webb, F.; Turner-McGrievy, G.; McKeever, B.W.; McKeever, R. Identifying and analyzing health-related themes in disinformation shared by conservative and liberal Russian trolls on twitter. *International journal of environmental research and public health* **2021**, *18*, 2159.
17. 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* **2021**, *118*, e2023301118, [<https://www.pnas.org/doi/pdf/10.1073/pnas.2023301118>]. doi:10.1073/pnas.2023301118.
18. Donkers, T.; Ziegler, J. The Dual Echo Chamber: Modeling Social Media Polarization for Interventional Recommending. Proceedings of the 15th ACM Conference on Recommender Systems; Association for Computing Machinery: New York, NY, USA, 2021; RecSys '21, p. 12–22. doi:10.1145/3460231.3474261.
19. Ghanem, B.; Buscaldi, D.; Rosso, P. TexTrolls: identifying Russian trolls on Twitter from a textual perspective. *arXiv preprint arXiv:1910.01340* **2019**.
20. Kim, D.; Graham, T.; Wan, Z.; Rizoio, M.A. Analysing user identity via time-sensitive semantic edit distance (t-SED): a case study of Russian trolls on Twitter. *Journal of Computational Social Science* **2019**, *2*, 331–351.
21. Im, J.; Chandrasekharan, E.; Sargent, J.; Lighthammer, P.; Denby, T.; Bhargava, A.; Hemphill, L.; Jurgens, D.; Gilbert, E. Still out there: Modeling and identifying russian troll accounts on twitter. 12th ACM conference on web Science, 2020, pp. 1–10.



22. Zannettou, S.; Caulfield, T.; Bradlyn, B.; De Cristofaro, E.; Stringhini, G.; Blackburn, J. Characterizing the use of images in state-sponsored information warfare operations by russian trolls on twitter. *Proceedings of the International AAAI Conference on Web and Social Media*, 2020, Vol. 14, pp. 774–785.
23. Wang, Y.; Callan, J.; Zheng, B. Should we use the sample? Analyzing datasets sampled from Twitter's stream API. *ACM Transactions on the Web (TWEB)* **2015**, 9, 1–23.
24. Bedi, P.; Sharma, C. Community detection in social networks. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **2016**, 6, 115–135.
25. Landherr, A.; Friedl, B.; Heidemann, J. A critical review of centrality measures in social networks. *Wirtschaftsinformatik* **2010**, 52, 367–382.
26. Das, K.; Samanta, S.; Pal, M. Study on centrality measures in social networks: a survey. *Social network analysis and mining* **2018**, 8, 1–11.
27. Moody, J.; McFarland, D.; Bender-deMoll, S. Dynamic network visualization. *American journal of sociology* **2005**, 110, 1206–1241.
28. Jacomy, M.; Venturini, T.; Heymann, S.; Bastian, M. ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software. *PloS one* **2014**, 9, e98679.
29. Roesslein, J. Tweepy: Twitter for Python! URL: <https://github.com/tweepy/tweepy> **2020**.
30. Chodorow, K.; Dirolf, M. *MongoDB - The Definitive Guide: Powerful and Scalable Data Storage.*; O'Reilly, 2010; pp. I–XVII, 1–193.
31. Kratzke, N. The #BTW17 Twitter Dataset - Recorded Tweets of the Federal Election Campaigns of 2017 for the 19th German Bundestag. *Data* **2017**, 2. doi:10.3390/data2040034.
32. McKinney, W. Data structures for statistical computing in python. *Proceedings of the 9th Python in Science Conference*. Austin, TX, 2010, Vol. 445, pp. 51–56.
33. Zaharia, M.; Xin, R.S.; Wendell, P.; Das, T.; Armbrust, M.; Dave, A.; Meng, X.; Rosen, J.; Venkataraman, S.; Franklin, M.J.; Ghodsi, A.; Gonzalez, J.; Shenker, S.; Stoica, I. Apache Spark: A Unified Engine for Big Data Processing. *Communications of the ACM* **2016**, 59, 56–65. doi:10.1145/2934664.
34. Kim, E.; Lee, K.; Yoo, C. On the Resource Management of Kubernetes. *2021 International Conference on Information Networking (ICOIN)*, 2021, pp. 154–158. doi:10.1109/ICOIN50884.2021.9333977.
35. Hagberg, A.A.; Schult, D.A.; Swart, P.J. Exploring Network Structure, Dynamics, and Function using NetworkX. *Proceedings of the 7th Python in Science Conference*; Varoquaux, G.; Vaught, T.; Millman, J., Eds.; , 2008; pp. 11 – 15.
36. Kernighan, B.W.; Lin, S. An Efficient Heuristic Procedure for Partitioning Graphs. *Bell System Technical Journal* **1970**, 49, 291–307, [<https://onlinelibrary.wiley.com/doi/pdf/10.1002/j.1538-7305.1970.tb01770.x>]. doi:<https://doi.org/10.1002/j.1538-7305.1970.tb01770.x>.
37. Parés, F.; Garcia-Gasulla, D.; Vilalta, A.; Moreno, J.; Ayguadé, E.; Labarta, J.; Cortés, U.; Suzumura, T. Fluid Communities: A Competitive, Scalable and Diverse Community Detection Algorithm, 2017. doi:10.48550/ARXIV.1703.09307.
38. Clauset, A.; Newman, M.E.J.; Moore, C. Finding community structure in very large networks. *Physical review. E, Statistical, nonlinear, and soft matter physics* **2004**, 70 6 Pt 2.
39. Langville, A.N.; Meyer, C.D. A Survey of Eigenvector Methods for Web Information Retrieval. *SIAM Rev.* **2005**, 47, 135–161.
40. Kleinberg, J.M. Authoritative Sources in a Hyperlinked Environment. *J. ACM* **1999**, 46, 604–632. doi:10.1145/324133.324140.
41. Aro, J. *Putin's Trolls: On the Frontlines of Russia's Information War Against the World*; Blackstone Publishing, 2022.
42. Kratzke, N. Monthly Samples of German Tweets, 2023. doi:10.5281/zenodo.7528718.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.