Online Social Networks and Media 31 (2022) 100228

Contents lists available at ScienceDirect

Online Social Networks and Media

journal homepage: www.elsevier.com/locate/osnem

Dissecting chirping patterns of invasive Tweeter flocks in the German Twitter forest

Jan Ludwig Reubold [a,*], Stephan Escher [b,*], Johannes Pflugmacher [b], Thorsten Strufe [a,b]

[a] Karlsruhe Institute For Technology, Am Fasanengarten 5, Karlsruhe, 76131, Baden Württemberg, Germany
[b] Technical University of Dresden, Nöthnitzer Str. 46, Dresden, 01187, Saxony, Germany

ARTICLE INFO

Keywords:
Social network analysis
Echo chambers
Community detection
User understanding

ABSTRACT

Twitter as a platform is used for news dissemination, with high volumes of campaigning and populism. This situation coincides with the growth of audiences who embrace social media as their primary news source. In general, effects like the deterioration of political education, misinformation, or ideological segregation then arguably represent a tremendous risk for democratic societies.

We analyze a comprehensive data set of the German-speaking Twitter community – a concise, well-defined Twitter population – to understand the extent and form of consumption of controversial news.

Our results affirm a high interest of German Twitter users in daily news and corresponding discussions. In-depth studies on the behavior, including tweeting- and grouping patterns, revealed the emergence of a new, more self-assured form of echo chambers.

1. Introduction

Online social networks, such as Facebook, Instagram, YouTube, and Twitter, attract enormous attention. These networks have almost ubiquitous reach. The information circulating in these networks is manifold and comes from various sources. In particular, news providers are making great efforts to publish and disseminate their articles on multiple social platforms to reach a wider audience [1]. Politicians, too, are embracing the digital environment. They use social media for campaigning and connecting with their target audience [2]. The amount and availability of informative content have caused a rising number of social media users to consume their daily news directly on these platforms [3]. The availability of social-media mobile apps amplifies this effect and increases exposure in various everyday situations. Democratized information acquisition, dissemination, and the free flow of information are positive aspects of this development. However, at the same time, it represents potential risks to political discussion in our society.

Additional actors have emerged. Some are distributing misinformation, conspiracy theories, and propaganda with agendas ranging from the commercialization of click-bait, over political influence, to establishing opinion platforms as hidden distribution channels for marketing of all types of products [4]. Recent publications underline that social media users are more exposed to populism, propagated by political actors from the extreme ends of the political spectrum, than individuals without social media [5]. A balanced news selection has to give way to a choice of posts and topics reinforced by the user's chosen neighborhood in this sheer mass of information. This fosters political polarization and ideological segregation [6–8]. Incidental news consumption reinforces such effects [9], leading to a reduction in political education [10]. This development arguably represents a primary risk for democratic societies. Individuals who put more trust in information shared by friends, likely regress to consume news from narrow contexts [9]. This development increases the difficulty of evaluating the credibility of information sources [11]. It consequently makes the emerging closed user groups more vulnerable to profit-oriented marketing, political campaigning, and general misinformation. Literature has termed such user groups *"echo chambers"*. A phenomenon that amplifies and reinforces common opinions within groups through repetition and mutual approval. Typically claimed to exist in social networks, they increase political polarization and ideological segregation [6]. Members of these *echo chambers* are more exposed to populism, propagated by actors from the extreme ends of the political spectrum [5]. They hence have a tremendous impact on the process of political opinion-forming.

Our goal is to analyze the impact of anti-democratic/con-troversial content on a western European Twitter community. In exhaustive studies, we try to answer the following research questions:

* Corresponding authors.
E-mail addresses: jan.reubold@kit.edu (J.L. Reubold), stephan.escher@tu-dresden.de (S. Escher), johannes.pflugmacher@web.de (J. Pflugmacher), thorsten.strufe@kit.edu (T. Strufe).
URL: https://ps.tm.kit.edu (T. Strufe).

https://doi.org/10.1016/j.osnem.2022.100228
Received 2 March 2022; Received in revised form 8 July 2022; Accepted 14 July 2022
Available online 30 July 2022
2468-6964/

- **RQ 1:** What is the extent and impact of controversial news content?
- **RQ 2:** To what extent are echo chambers feeding on controversial content influencing the community?

Due to data collection limitations, we have to strike a trade-off between sample size and data quality. We base our studies on a concise, well-defined, and virtually complete Twitter community, concentrating on the German-speaking Twitter community (GTC).

Selecting tweets by the language allows for a detailed observation of such a specified population. Often, culturally and geographically diverse groups speak the same language. The GTC, however, represents a large, geographically well-defined population of around 7 million active users. The majority are from Germany, Austria, and adjacent parts of neighboring countries, and they all share a relatively homogeneous political landscape and corresponding media outlets.

To study anti-democratic content, we define controversial and non-controversial content. Controversial content combines articles from providers that contribute to misinformation, conspiracy theories, political propaganda, and similar democracy decomposing elements.

We base our studies upon two building blocks:

- **Content** We propose an automated data augmentation strategy to facilitate data enrichment on large, real-world data sets. We leverage shared external content and hashtags to get a high-level understanding of discussions in an automated manner.
- **Distribution/Impact** We leverage dynamic interactions between users (i.e., mentions, retweets, quotes, and replies) to accurately measure relationships.

Thereby, we get a (*i*) high-level understanding of what is shared/discussed based on the automated categorization of content, a (*ii*) measure on the share of news related discussions within the network, and (*iii*) can identify influential actors and multiplication networks, (*iv*) measure the presence of established news providers within the network, (*v*) analyze user engagement w.r.t. different types of news, (*vi*) study discussion patterns of users, and (*vii*) perform a community structure analysis.

Based on a comprehensive understanding of the content contributing to political opinion-forming, we study the influence of phenomenons related to controversial content.

In the remainder of this paper, we first give an overview of the state-of-the-art in Section 2 and describe our approach in Section 3. We report on the results of our experiments and discuss them in Sections 4 and 5, to conclude with a summary in Section 7.

2. Mensuration of news consumption and identification of community structures

Online social networks (OSNs) offer researchers the opportunity to investigate human interaction on a large-scale [12,13]. Based on such rich data, we address various topics of technical and theoretical nature. In the following, we provide a brief overview of current research on topics related to our approach.

2.1. Political orientation

Studies on the matter of OSNs rely on large and rich data sets. Depending on the objectives, data often has to be augmented with further information. Our studies rely on information about the political affiliation of users.

In this context, Colleoni et al. [8] worked on the complete Twittersphere of 2009 provided by Kwak et al. [14] to investigate the political homophily of Republicans and Democrats across the entire network. Using linguistic features extracted from annotated tweets and news texts, they utilized a supervised classification approach. While a common approach in the area of user classification on Twitter [15], research showed that the prediction of political affiliation is not reliable in multi-class scenarios, e.g., in the context of the broader political spectrum of German parties [16]. Additionally, textual features of tweets are not stable over time. Here, tweets from topical authorities, who seem to be more consistent in their messages, represent an exception [15].

Therefore, an algorithm inferring user characteristics and interest from context-specific activities is more promising for the German Twitter user base. In this context, several attempts rely on Wikipedia articles to infer the interests of users [17–19]. Wikipedia and its broad range of categorized articles, including people, events, and locations, can be utilized to build a reliable knowledge database. Faralli et al. [20] approximated user interests by finding *"friends"* they could link to Wikipedia articles. For example, if a user followed a famous basketball player, her interests included sports and basketball. The researchers proposed a hierarchical representation of user interests and conducted a large-scale homophily analysis on Twitter. Their methodology offered a compact, tunable and readable way to examine user interests.

For a more thorough understanding of user interests, Himelboim et al. [21] leveraged frequently shared hyperlinks, user mentions, and hashtags and, thereby, analyzed users based on domain-related interests and hashtags. We deploy a similar approach for inferring user attributes.

2.2. Community detection and echo chambers

We explore the existence and spread of echo chambers. Detecting echo chambers is commonly performed by modeling the network as some graph and extracting clusters of nodes with high interconnectivity. Therefore, we rely on the information of the community structures within our data. Besides the investigations on echo chambers, we also leverage the information to understand user behavior.

Early studies investigated simple social graphs, as represented in the contact relationships on the OSN [22,23]. These approaches assumed that online friendships inherit crucial attributes from real-world relations so that the majority of meaningful interactions in OSNs occur between friends. Wilson et al. [24] analyzed interactions, i.e., wall posts and photo comments, among Facebook users. They reported that most user interactions occur only within a tiny subset of a user's friendslist, often leaving half of the remaining friends out of all communications. They studied an interaction graph containing only edges between users interacting instead of relying on friendship links. Their evaluations on two adequate social applications demonstrated that using an interaction graph yields better results than using a friendship graph.

Himelboim et al. [21] used topic networks applied to a clustering approach to detect echo chambers on Twitter. They collected topic-related Twitter data and created multiple interaction graphs based on retweet-, mention-, and reply relationships. Using the Clauset–Newman–Moore algorithm, they identified communities of users that had frequently interacted with each other. These users created a structure of interaction silos where echo chambers might emerge. They then assessed the occurrence of ideological similarities among users within a community by analyzing their frequently shared hyperlinks, user mentions, and hashtags for a more thorough examination of the identified communities. By assigning a political affiliation to influential users within the community, they aimed to infer its political orientation. An influential user was determined with in-degree centrality to measure his exposure to other users.

Conover et al. [25] demonstrated that detecting exposure to alternative news vs. segregation into echo chambers yields different, possibly conflicting results depending on the chosen interaction to model the graph. Their experiments illustrate the importance of the selected methods and graph modeling schemes. Utilizing two different interaction graphs, they tried to link user similarities and political orientation. Crawling 250 000 tweets during the 2010 U.S. congressional midterm elections, they modeled graphs both, depending on retweets and mentions. Detecting communities on the Retweet-graph resulted in two highly segregated communities with opposing political ideologies and only a few inter-connections. The authors concluded that the structure encoded user preferences to retweet similar political views. Community detection on the Mention-graph, however, yielded fundamentally different results. A single community emerged, containing politically heterogeneous users. Despite their opposing political ideologies, these users exhibited a high level of interaction. They concluded that users from both political sides confront each other with content, contrary to their political affiliation, which leads to a ruffled exchange of Tweets. In conclusion, they posit that meaningful analysis benefits from comprehensive, combined interaction graphs.

Besides Conover et al. several other studies in recent years dissected echo chambers within OSN. These reports describe a user's tendency to retweet content with political views similar to theirs [25–27]. They observe sharing content from such a narrow context fosters segregation by political orientation.

Other studies on this topic investigate the extent of such behavior considering the political orientation of individuals [6,8]. Boutyline and Willer [7] observed that conservative and politically more extreme individuals showed a more pronounced tendency to form segregated user groups than liberals. While Barberá et al. [6] report similar results consistent with psychological theory and research bearing on ideological differences in epistemic, existential, and relational motivation, they conclude that previous work may have overestimated the degree of ideological segregation in online social networks.

In this context, we have to emphasize the difference between the concept of an echo chamber and an epistemic bubble. While the latter relates to information networks that exclude important information sources without their members noticing it, echo chambers actively discredit or even exclude contrary opinions [28].

2.3. Promotional profiles

Besides manually controlled accounts, there also exist orchestrated and automated ones. Several guidelines recommend creating social media profiles for improved public relations and dissemination. To increase the distribution of news content in social media, Orellana-Rodriguez et al. [29] propose best practices. They suggest creating employee accounts to promote their corresponding tweets. Such accounts should contain a statement about their affiliations. News providers establish Twitter profiles to further the distribution of their articles [1].

News agencies, such as Reuters or AFP, instruct journalists using their accounts for work to include a disclaimer. The disclaimer identifies them as employees of a specific news agency [30–32]. It should also include a declaration that they speak for themselves and not their employers.

3. Dissecting German tweeting flocks

This work provides exhaustive studies on the news consumption of German-speaking Twitter users. The basis of our approach is the data acquisition strategy (i.e., obtaining an automatically labeled, virtually complete data set) and the sophisticated, improved modeling of interaction graphs. We assume that measurements of the shared external content allow us to approximate statistics on news consumption. The classification into categories allows for an automated high-level understanding of its content. Additionally, hashtags (related to shared external content) provide further semantic understanding. The approach avoids biases due to inaccuracy during the pre-processing. An example here is utilizing NLP techniques for semantic understanding.

In the following, we introduce the various parts of the data engineering process. Therefore, we summarize Twitter functionalities before presenting our data collection strategy and explaining the automated data enrichment (e.g., promotional profile detection, domain categorization). We conclude with a detailed discussion on sophisticated interaction graph modeling.

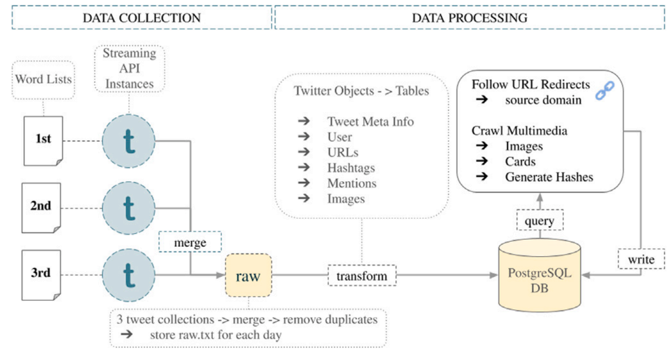

Fig. 1. Data collection pipeline with three parallel Twitter Streaming API instances; each with a separate stop word list, including 400 frequently used terms in the German language; output of streams is merged; duplicated entries are dropped; raw Twitter-Objects are extracted from the files and parsed into a PostgreSQL database.

3.1. Twitter OSN and functionalities

Twitter offers its users different types of *Tweet-Objects* to generate content on the platform. As of 2019, a user can write a message to his timeline, also known as a status update. The timeline of a user represents a roster of posts. It records activities and makes them visible to followers. The *following functionality* represents the core of the Twitter eco-system. Based on the accounts a user follows (e.g., news providers, celebrities, and friends), Twitter compiles an overview of current events and activities. This feed displays activities of *followed* others to whom the user has subscribed. Therefore, the system provides a news-feed-like overview tailored to the user's choice.

On Twitter, the *original tweets* is the standard way of posting. *Retweets* represent another type of post, which allows a user to copy a tweet from another user to his timeline. Therefore, it is visible to his respective followers and visitors. Users can also *quote* other tweets (except retweets). Thereby, they can re-post a user's message with a comment of their own. Lastly, there are *replies* to comment on any given tweet, except retweets.

Besides textual content, such a *Tweet-Object* can also contain multimedia content (*photos*, *videos*, *animated GIFs*), interactive content (*hashtags*, *user mentions*), places (*geolocation*), or links (*URLs* linking to external sources, which commonly are visualized as *Twitter Cards*). In addition to manually embedded user mentions (@username), Twitter automatically adds *mentions* in front of content that implies an interaction between users (retweets, replies, and quotes). Further, every Tweet-Object has an attribute (*source*) that describes the service used to post the tweet. We extracted the service from each tweet in our data set to estimate their usage. Besides official Twitter clients, there are also third-party services. These services allow accounts to post tweets in an automated manner.

User-objects provide a variety of meta-data. It contains multiple free-text fields (e.g., name, description, URL), statistics about the social links of a user (e.g., follower-, and friend count), and statistics about her activities (e.g., favorites-, and tweet count).

Users can interact with others via direct User Mentions within a tweet or indirectly via connected tweet types, such as retweets, quotes, and replies. Compared to static follows, interactions allow capturing relationship dynamics over time.

3.2. Data acquisition

The studies depend on a virtually complete snapshot of our target community. Therefore, we propose a comprehensive data collection scheme, i.e. an extension of Scheffler's approach [33] (cmp. Fig. 1).

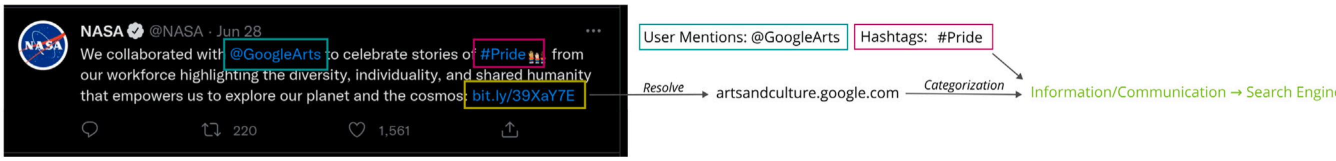

Fig. 2. Identification of meta-information in large-scale networks; Information at the core of our research: *User Mentions*, **categorized** *Hashtags*, and **categorized** *URLs*.

Our evaluation of different collection methods confirmed Scheffler's findings. Geolocation-based filters only capture tiny amounts of German tweets. We hence decided to utilize word lists for our purpose. In contrast to Scheffler, we do not collect-then-filter to remove tweets in other languages, but we leverage the built-in language identification of Twitter. We thus created word filters, encompassing the 1200 most frequent German words. We base our choice on multiple text corpora, provided by the Leipzig Corpora Collection [34] and one corpus of frequently used words from OpenSubtitles.org.[1] The latter encompasses terms that are more prevalent in informal conversations. Twitter enforces a maximum of 400 keywords per instance, so we divided our word filter into 3 different lists and used three individual, parallel data streams. All streams obtained many tweets from 600k to 1.2M on average. Thus our approach does not exceed the rate limitations of 1% ($\approx$ 5M tweets). We drop duplicated entries and merge the stream outputs.

Findings in Morstatter et al. [35] suggest that German tweets are sufficient to capture political debates of the German-speaking population as non-German Tweets are ignored by the community. So, relying on Twitter's language detector, we exclusively capture German tweets. Therefore, we sidestep Twitter's rate limitations and, thereby, avoid down-sampling. While the detector lacks thorough documentation, research showed that, in some cases, it outperforms established alternatives such as Google's Compact Language Detector [36].

We enrich recorded tweets with additional data. Besides the attributes, we further extracted child objects (original tweets, replies, quotes) from collected Tweet-Objects. The latter may entail collecting additional (non-German) Tweet-/ User-Objects. We argue that we need to include users who do not tweet in German but interact with German tweets.

3.3. Data enrichment

We obtained a virtually complete snapshot of the GTC, collected during 2 months surrounding the European Parliament Election in 2019 [37]. Still, we have to understand the content to analyze news consumption on Twitter. We base subsequent studies on this understanding. Thus, we need a robust, generalized strategy. In the following, we propose an automated, sophisticated, and comprehensive data enrichment strategy evolving around shared external content (see Fig. 2).

We focus on embedded news: shared external links presented as a preview within the social media platforms (for instance, Twitter cards with a headline, thumbnail, and summary on Twitter). Our analysis terminally requires to extract the category and type from the shared tweets as well as additional meta-information (e.g., promotional profile, controversial user), which we perform in the following ways:

3.3.1. Understanding content

We begin by augmenting tweets with meta information to obtain a high-level understanding of their contents.

Functional groups: Categorizing domains. We categorize domains leveraging McAfee's TrustedSource[2] (2019) to obtain a comprehensive understanding of the shared external content. We tested different categorization services, and McAfee's TrustedSource successfully identified the highest number of domains. Further, it provides a fine-grained set of 100 hierarchical categories (e.g., News, Lifestyle, Political Opinions, or Spam). McAfee also provides semantic subsets that split the categories in 12 so-called Functional Groups (FGs). Using TrustedSource, we categorized 98.3% of the URL-tweets in our data set. In the remainder, we sort URLs based on their domains into **FGs** and its related **categories** (**FG→category**).

News group. To identify all domains that influence the forming of political opinions, we manually investigated the most-shared websites from every category in our data set. Based on this research, the following set of domain categories, distinguished by the objectivity of reports (from moderate-, over tendentious- to extreme views), comprises the ***News Group***:

- Information/Communication → **General News:** Domains that generate daily news, political opinion sections, and educational content.
  Top 5: spiegel.de, welt.de, bild.de, sueddeutsche.de, and *zeit.de.*
- Society/Education/Religion → **Education/Reference:** Web pages that relate to educational content, for example, classic literature, history, art, and other academic-related content.
  Top 5: de.wikipedia.org, spektrum.de, fridaysforfuture.de, kurierdeswissens.de, danisch.de.
- Society/Edu./Religion → **Non-Profit/Advocacy/NGO:** Web pages run by charities and or educational groups or campaigns.
  Top 5: change.org, correctiv.org, peta.de, deutschland-kurier.org, mimikama.at.
- Society/Education/Religion → **Government/Military:** Web pages provided by governmental or military organizations, including national branches as well as supranational entities, such as the United Nations or the European Union.
  Top 5: bundestag.de, polizei.bayern.de, auswaertiges-amt.de, bundeswahlleiter.de.
- Society/Education/Religion → **Major Global Religions:** Web pages that provide information about major religions (e.g., Buddhism, Chinese Traditional, Christianity, Hinduism, Islam, Judaism, etc.) and include discussions and non-controversial commentary.
  Top 5: katholisch.de, catholicnewsagency.com, kath.net, vaticannews.va, evangelisch.de.
- Society/Edu./Religion → **Politics/Opinion:** Web pages that cover political parties and opinions on various topics such as political debates.
  Top 5: tichyseinblick.com, jungefreiheit.de, achgut.com, politikstube.com, volksverpetzer.de.
- Lifestyle → **Controversial Opinions:** Web pages that share extreme opinions, which are offensive to political or social sensibilities. Examples include xenophobic, fundamentalist viewpoints, and disinformation campaigns.

[1] https://github.com/hermitdave/FrequencyWords/

[2] https://trustedsource.org

Top 5: journalistenwatch.com, pi-news.net, philosophia-perennis.com, anonymousnews.ru, der-dritte-weg.info.

- Risk/Fraud/Crime → **Discrimination:** Web pages that provide content that explicitly encourages the oppression or discrimination of a specific group of individuals. There are only a few domains that McAfee classifies as discrimination and only a few found in our data.
  Top 5: metapedia.org, theeuroprobe.org, renegadetribune.com, vanguardnewsnetwork.com, nordfront.se.
- Risk/Fraud/Crime → **Historical Revisionism:** Web pages that spread misinformation, or offer divergent interpretations of, significant historical facts (e.g., Holocaust denial).
  Top 5: renegadetribune.com, vho.org, altright.com, dailystormer.name, johndenugent.com, codoh.com.

Accessing OSN links. In addition to external sources referring to news content, we want to gain insight into content included in links to posts on other social platforms. Thus, we developed web crawlers for *YouTube*, *Facebook*, and *Instagram*, the most shared platforms in our data set. By utilizing the *YouTube Data API v3* and the *HTML* and *JavaScript* sources from Facebook and Instagram, we identify corresponding external profiles and their influence on news distribution on Twitter.

Political hashtags. To investigate user discussions about shared news content, we also consider the hashtags they contain. We automatically categorize the corresponding tweets by leveraging the co-occurrence of hashtags with URLs, as classified above. For example, if the hashtag #CDU appears in a tweet that also shares an article from *Spiegel*, we assign the #CDU hashtag to the category General News. This approach allowed us to assign categories to 60% of the hashtags in our data set. Note, however, a hashtag is assigned to several categories depending on its usage w.r.t. URL-tweets.

3.3.2. Understanding effect

Besides understanding content, we also want to study its distribution and impact. Therefore, we introduce the following definitions.

User engagement. To measure how *users engage* with news or external political content, we define *reaction-tweets* in addition to simple tweeting and retweeting. Reaction-tweets contain direct responses (replies and quotes) and their retweets. We attribute them to the original tweets they are referencing.

To measure content popularity, we leverage related reactions. We also identify self-promotional profiles and influential users to complement our studies on *engagement*.

Promotion profiles and automated accounts. We identify promotional profiles to measure their impact on the distribution of news. We base our automated detection of self-promotional profiles on the guidelines of news agencies such as Reuters or AFP (see Section 2.3). This process yields two types of promotional profiles: (*i*) journalists and (*ii*) feeds (see Table 1). We identify a journalist's profile by checking if it stated a news source in the free-text *URL-field* (e.g. *spiegel.de*) as well as the respective news domain in the *description text* (e.g. *Spiegel*) of their profile. Feeds, we identify using the above and check if their screen name contains the respective news domain (e.g., @spiegelonline).

These feeds often act as the official publisher of articles. They generate automated content and rarely interact with other users. Websites often create multiple feeds solely to disseminate their articles.

This approach has obvious limitations. We cannot automatically detect promotion profiles that do not follow the journalistic guidelines. Therefore, we conducted a manual search for additional promotion profiles for the 30 top content providers. As it did not yield any additional profiles, we are confident that our findings below are representative in this regard.

Besides official automated profiles, malicious bots exist. With regard to these bots, we pursue a different route. In general, bot detection is an unsolved problem. For this reason, scientists resort to heuristics. Often, suspended accounts are interpreted as bots. However, a recent study [38] reports that less than 1% of the suspended accounts were suspected or potential bots. In line with other research, they found that suspended accounts pursued specific polarizing political agendas. Another approach to identify bots is to use tools such as the BotOrNot service. While often used by scientists, research shows how limited this approach is [39,40]. With Twitter adding that binary judgments have real potential to poison our public discourse.[3] Based on this evidence, we argue that using these heuristics to exclude bots from our study provides no guaranteed benefits while seemingly introducing significant amounts of noise.

Table 1
Sample of self-promotional Twitter-profiles from Spiegel.

Screen name	@SPIEGEL_Politik	@joleffers
Name	SPIEGEL ONLINE Politik	Jochen Leffers
Description	Hier twittert das Politik-Ressort von @SPIEGELONLINE. Datenschutz: http://spon.de/afemu	ist bei SPIEGEL ONLINE im einestages-Ressort, twittert hier aber-so-was-von-privat
URL	spiegel.de	spiegel.de
Journalist	✗	✓
Feed	✓	✗

Influential users. Researchers proposed different measurements to identify influential accounts on Twitter [41]. In this work, we follow the approach of Kwak et al. [14] by applying the PageRank algorithm to our network. However, we augmented the approach in two ways. Instead of the passive topology metric (i.e., follower-links) – a poor indicator of actual influence [42] – we utilize interaction activities of users (i.e., retweets, mentions, quotes, and replies) to form our edges. Therefore, our approach relies on similar information as Himelboim et al. [21] (see Section 2). However, we do not rely on an undirected network, assigning symmetric values to interactions between users, but construct a directed graph by calculating scores that indicate how much a user interacts with another user. The resulting weighted PageRank score for each user contributes to a more precise examination of influential nodes in our network.

Furthermore, we expand our research regarding the detection of influential news providers. We determine the influence of news providers by influence measurements. These influence measurements include provider abilities to spread news articles and how many users they can reach. Our approach complements usual methods to measure the popularity profiles in online social networks (e.g., surveys) [3]. In contrast to surveys, a methodology based on the sharing and commenting on news provides a more detailed depiction of user behavior. Also, unlike surveys based on self-reports, it is not vulnerable to social desirability bias [43]. PageRank measures the global influence of nodes in a network and, thereby, lends itself to this task.

3.3.3. Understanding users

So far, our data enrichment strategy allows us to understand the content and distribution of tweets. However, we also want to gain insights into the political attitude of users. Therefore, we augment user information by leveraging their interests.

User interests. Prior work [20] identified user interests based on language processing and augmented this information into the friendship graph. This approach yields a more static assignment and relies on potentially error-prone text extraction. We aim to capture the dynamics of interest more accurately. Therefore, we identify it according to the

[3] https://blog.twitter.com/en_us/topics/company/2020/bot-or-not

hyperlinks the users interact with and share to avoid language processing and ambiguities. Using our approach, we leverage the categories of shared domains and hashtags. Briefly, we consider a user who regularly shares or replies to a specific news domain interested in related topics.

Controversial users. The majority of studies classify users based on a political spectrum. Expressing opposing views in the political landscape of the U.S., researchers often label users as either Democrats or Republicans. Since the U.S. has a virtually two-party system, this is a justified and sensible approach. The political landscape in German-speaking countries, however, is more diverse. The political agendas of parties, e.g., tend to overlap. Also, deducing opinions based solely on hashtag information does not distinguish between support and opposition. Therefore, we do not rely on party references in tweets for estimating political affiliation.

The 'Hidden Tribes' study [44] took a more nuanced approach to analyze America's political landscape. Surveying 8000 Americans, they identified seven groups based on shared beliefs and behaviors. Interestingly, the groups furthest to the right and left of the political spectrum were similar in surprising ways (e.g., color and wealth) and, most importantly, these two groups are the driving force behind the widening of the gulf between the two political factions. Therefore, we distinguish between moderate and extreme users, labeled as non-controversial and controversial.

Based on McAfee's TrustedSource database, our domain categorization approach identifies domains that produce extreme political content and misinformation. While the category *Politics/Opinion* already contains domains with extreme and inflammatory content, categorizing users as controversial based on a shared article of these domains would lead to imprecise labels. Hence, we only include domains with extreme political views that, e.g., deny the Holocaust or encourage the oppression or discrimination of specific groups.

In this context, we assume that retweeting indicates an interest in a topic or even agreement with the sentiment of a message [45]. Therefore, after investigating all of the domains, we posit that people, who support these contents by sharing them in the network and contributing to its distribution, are likely to hold extreme political views. The categorization in our database classify these domains as Controversial Opinions, Discrimination, and Historical Revisionism. We define a group of **Controversial Users** comprised of users that shared at least one of these URLs. Accordingly, we specify users who share non-controversial content as **Non-Controversial Users**. While we cannot deduce their political affiliations, we assume they manifest less extreme views.

3.4. Extracting interaction graphs

Our strategy for data acquisition- and augmentation provides an understanding of tweet content and distribution. However, further information on interactions is necessary to understand news-related dynamics. Therefore, we introduce a sophisticated modeling scheme for interaction graphs.

Two major approaches exist, where one utilizes static follower-relations, the other leverages dynamic interactions between the users in the network. The *following*-functionality of Twitter offers users a way to keep track of each other's content. By following the activities of others, users express endorsement or even take part in sweepstakes. Since links between users can be one-sided or reciprocal, many users try to expand their influence in the network by offering reciprocal following-relationships to like-minded people. However, follower-relations are insufficient to understand the relationships [24].

Interactions between users can either be found in direct user mentions or indirectly by using connected tweet variants, such as retweets, quotes, and replies. In contrast to static follower-relations, these interactions gather more information about relationship dynamics over time. For example, users frequently retweeting each other's content during a political election seem to share the same political orientation. Retweeting indicates that a user is interested in a topic or even agrees with the sentiment of a message [45]. An extensive (reciprocal) retweeting among users could also show a certain level of trust and appreciation for each other. Quotes allow to retweet content with additional commentary. It can either express opposition or praise to the quoted post and its originator. The use of mentions and replies is more prevalent between users having opposing views on a specific topic [25]. While retweets provide no platform for further interaction, quotes and replies allow for comment. Thus, reaction-tweets can start discussions.

Examining conversations within Twitter is a promising strategy to gain insight into the relationship between users. Therefore, we rely on user interactions. In the following, we describe the community detection algorithm.

3.4.1. Interaction graph

We want to model the exposure of users and their communities to news and categories. For that purpose, we model the users V as the vertices and all Twitter interactions as connecting weighted edges within a graph. We want the weights to represent similarity for later community detection.

The semantics of distance in social graphs depends on the type of interaction. Gadek et al. [46] posit that quantified interaction is a promising metric to estimate a distance between users. We thus quantify interactions between users, combining the four interaction types: retweets, replies, quotes, and user mentions. Each interaction has its own semantic. Therefore, we calculate one metric for each interaction type and accumulate the scores to a final edge weight, denoting the distance.

Given a set of N users, $U \triangleq \{u_i\}_{i=1}^{N}$, and the different types of tweets Ω,

$$\Omega \triangleq \{\alpha = \text{original tweet}, \beta = \text{retweet}, \gamma = \text{reply}, \tau = \text{quote}\},$$

we break down the count of all tweets T by their type with $T_\omega\left(u_a, u_b\right)$ as the total number of tweets of $\omega \in \Omega$ user A posted. When A posts an original tweet, B is the empty set $\emptyset$ Otherwise, B is the author of the original tweet. Further, the total number of tweets T user A posted, for example, is expressed as

$$T_\Omega\left(u_A, \cdot\right) = \sum_{\omega}^{\Omega} \sum_{u}^{U} T_\omega\left(u_A, u\right).$$

Accordingly, $T_{\Omega\setminus\beta}\left(u_A, \cdot\right)$ represents the total number of tweets of user A that were not retweets.

The Retweet score is based on (i) the number of retweets from tweets of user B shared by user A ($T_\beta\left(u_a, u_b\right)$) and (ii) the number of all tweets of users B that are no retweets ($T_{\Omega\setminus\beta}\left(u_b, \cdot\right)$ and defined as

$$S_\beta\left(u_a, u_b\right) \triangleq \frac{1}{2}\left(\frac{T_\beta\left(u_a, u_b\right)}{T_{\Omega\setminus\beta}\left(u_b, \cdot\right)} + \frac{T_\beta\left(u_b, u_a\right)}{T_{\Omega\setminus\beta}\left(u_a, \cdot\right)}\right). \quad (1)$$

Note that we exclude retweeted retweets from the equation because these tweets essentially are retweets of the original tweet, e.g., $A \rightarrow B \rightarrow C$ we capture as $A \rightarrow C$.

The idea behind the metric is that user A retweets a specific number of tweets from user B. The more content users retweet from each other, the closer their distance in the graph. For example, if A retweets every tweet from B, they are closer together in the graph since A shares the same content as B. Therefore, two profiles that were to retweet each other's every tweet, virtually mirroring one another, would represent the closest profile distance.

The corresponding scores for quotes and replies are defined accordingly, as

$$S_\tau\left(u_a, u_b\right) \triangleq \frac{1}{2}\left(\frac{T_\tau\left(u_a, u_b\right)}{T_{\Omega\setminus\beta}\left(u_b, \cdot\right)} + \frac{T_\tau\left(u_b, u_a\right)}{T_{\Omega\setminus\beta}\left(u_a, \cdot\right)}\right), \quad (2)$$

and

$$S_\gamma\left(u_a,u_b\right) \triangleq \frac{1}{2}\left(\frac{T_\gamma\left(u_a,u_b\right)}{T_{\Omega\setminus\beta}\left(u_b,\cdot\right)}+\frac{T_\gamma\left(u_b,u_a\right)}{T_{\Omega\setminus\beta}\left(u_a,\cdot\right)}\right). \tag{3}$$

In contrast to the other interactions, user mentions are not tweet-variants but interactive elements added to tweets. Every tweet potentially contains a User Mention that links a specific user profile. Profiles hence are closer to each other if they have frequent, mutual mentions. For calculating the User Mention metric, we need two statistics: The number of user mentions between respective users and the total number of user mentions per user. We then encode Mentions similar to tweets and define the User Mention Score as follows:

$$S_M\left(u_a,u_b\right) \triangleq \frac{1}{2}\left(\frac{T_M\left(u_a,u_b\right)}{T_M\left(u_b,\cdot\right)}+\frac{T_M\left(u_b,u_a\right)}{T_M\left(u_a,\cdot\right)}\right). \tag{4}$$

Our final *Interaction Score* combines all interaction metrics mentioned above as

$$S\left(u_a,u_b\right) \triangleq \frac{1}{4}\sum_{v}^{\tilde{\Omega}} S_v(u_a,u_b), \tag{5}$$

with $\tilde{\Omega} \triangleq \{\beta,\gamma,\tau,M\}$, representing the mean value of all scores combined.

S thus encompasses all interactions between users, and we apply it as the final weights to the edges of our graph. The edge weight ranges from 0 to 1. A higher score results in closer distances in the graph, therefore supporting the detection of user groups that frequently interact with each other. For studies on the communities, we perform additional analyses on the separate metric scores (1) - (4).

3.4.2. Community detection

We define a community as a sub-graph of the network. The literature distinguishes between soft- and hard clustering, where nodes may be associated with several different communities in the former, but only a single one in the latter case. Soft clustering commonly identifies much higher numbers of communities compared to hard clustering. For a fine-grained analysis, soft-clustering hence is intractable on such massive graphs. Further, focusing on the big picture of the network at hand, the results of a hard clustering approach identify communities that are most densely connected. Due to the large scale of the graph, we chose to apply the Louvain method [47]. It represents a hard-clustering approach based on a greedy algorithm that optimizes modularity. It runs for several iterations on weighted graphs and detects hierarchies of clusters in this process. The hierarchical partitioning of communities allows for a more detailed analysis of the discovered communities.

An issue with modularity optimization is the so-called resolution limit, i.e., there is no guarantee to detect small communities or combinations of small, weakly interconnected communities. The discovered structure does not necessarily correspond to the most pronounced community structure. Fortunate and Barthelemy studied the effects of the resolution limit and questioned the usefulness of modularity in practical applications [48]. However, by utilizing the hierarchical approach of the Louvain method, we can circumvent the issue. Taking successive iterations into account, we can identify small communities in early iterations of the hierarchical process.

3.4.3. Quality indicators

There is no correct quality assessment strategy to measure the goodness of fit of an identified community structure. However, related work leverages various indicators for this task. Besides the modularity score, the size distribution and the ratio between intra- (communications within a community) and inter-scores (communications between communities) serve as indicators.

Table 2
Twitter-Objects captured during the data collection process.

Object-Type	Count	Tweets (%)	Users (%)
Tweet	77 390 122	–	–
User	6 919 206	–	–
Mention	85 155 158	72	80
URL	18 358 074	23	25
Hashtag	39 197 019	22	29
Multimedia	19 702 261	19	56
Place	1 189 696	2	2

Table 3
Distribution of Tweet variants when performing actions.

Action	Tweet variant		
	Original (%)	Reply (%)	Quote (%)
Retweeting	66.8	21.7	11.5
Replying to	24.7	71.7	3.6
Quoting	76.5	14.5	9.0

Modularity. Modularity measures the difference between the original graph and a randomized graph. The value ranges between −1 and 1, where a positive value indicates that the edges within communities exceed the expected connectedness compared to random connectivity. According to Reichardt and Bornholdt [49], the expected maximum modularity for a random network is $Q = 0.15$. Wang [50] compared modularity values across different clustering algorithms. They reported that $Q \geq 0.4$ is a sufficient threshold for detecting meaningful, distinct communities in a graph.

Size distribution. The Interaction Graph represents a network built on the individual activities of people using Twitter. Therefore, the community structure potentially includes both small groups and large communities. A commonly observed indicator of real networks is the heterogeneity of their size distribution. It means most community detection methods find skewed distributions of community sizes [51–53].

Score-ratios. Based on the average interaction score, we compute the ratios. The desired outcome, a higher score for intra-edges, indicates communities with more densely connected users. On the other hand, a higher value for inter-edges suggests a slight imprecision in the separation of communities.

4. Analysis

In the following, we present our exhaustive studies on the news consumption characteristics of German-speaking Twitter users. Our goal is to answer questions on controversial news content (see RQs [1]). However, we require supplementary studies to classify findings regarding controversial users. Therefore, we examine different aspects related to news content.

In the following, we introduce the data set and report relevant statistics. We study news-related content and its overall share within the network leveraging figures on *Functional Groups* (URL categories for automated content understanding, see Section 3.3.1), hashtag-usage, and external OSN content. We explore distribution patterns, reach and impact. Then, we complement our studies by analyzing the behavior patterns of news-interested users. Finally, we turn to controversial users. Here, we concentrate on the behavior of controversial users in the Twitter community and investigate the existence of isolated controversial users groups.

4.1. The German-speaking Twitter community

To obtain a representative sample of the Twitter-sphere of the German-speaking user base, we collected tweets throughout two

Table 4
Most shared Hashtags during busiest days of data collection; here GTNM stands for Germany's Next Top Model.

Period	Hashtag	Count	Category	Event
May, 26th – 28th	#Europawahl2019	105 498	politics	European Parliament Election 2019
	#CDU	42 570	pol. party	European Parliament Election 2019
	#AfD	36 038	pol. party	European Parliament Election 2019
	#EUWahl19	25 562	election	European Parliament Election 2019
	#AKK	24 697	politician	European Parliament Election 2019
	#Europawahl	24 185	election	European Parliament Election 2019
	#Rezo	23 498	controversy	European Parliament Election 2019
	#SPD	21 519	pol. party	European Parliament Election 2019
	#Zensur	16 087	controversy	European Parliament Election 2019
	#Grüne, #Grünen	14 290	pol. party	European Parliament Election 2019
	#Sachsen	11 272	election	European Parliament Election 2019
June, 2nd	#NichtOhneMeinKopftuch	124 218	politics	Political campaign
	#Nahles	17 991	politician	Politician resignation
	#SPD	14 698	pol. party	Politician resignation
	#뷔	9 237	music	Campaign by BTS
	#태형	9 234	music	Campaign by BTS
	#태태	9 066	music	Campaign by BTS
May, 18th	#EurovisionSongContest2019	67 188	music	Eurovision Song Contest 2019
	#Eurovision	45 210	music	Eurovision Song Contest 2019
	#Strache	30 218	politician	Ibiza Affair
	#esc2019	23 038	music	Eurovision Song Contest 2019
	#strachevideo	12 911	controversy	Ibiza Affair
May, 22nd – 23rd	#rezo	30 742	controversy	European Parliament Election 2019
	#GNTM	22 980	entertainment	TV Show
	#CDU	20 216	pol. party	European Parliament Election 2019
	#Amthor	19 006	politician	European Parliament Election 2019
	#EuropaWahl2019	15 595	election	European Parliament Election 2019
	#RezoVideo	14 515	controversy	European Parliament Election 2019

months — between the 2nd of April and the 2nd of June 2019. A detailed discussion on the diversity of the content captured during that period can be found in Appendix C. The sample contains 77 million tweets and 6.9 million user profiles (ref. Table 2 for an overview).

4.1.1. Tweet types

Categorizing these Tweet-Objects by tweet type (i.e., original tweet, retweet, reply, quote) revealed that the most frequent action was retweeting. The majority of activity in our sample was reactive. Retweets account for 38% of all tweets in our corpus and are used to distribute content from other users, Replies for 31%, and original tweets, creating novel content or initiating conversations, account for only 27%. Quotes are rarely used at all (3.7% of the sample). Interestingly, we observed fewer users in our sample using replies (23%) than retweets (64%).

Besides investigating tweet types, we also analyzed their interactions. Table 3 shows that most often original content was retweeted (66.8%), followed by replies (21.7%) and quotes (11.5%). Looking at quoting, the distribution is very similar. Regarding replies, however, most of these tweets react to other replies (71.7%).

4.1.2. Tweet content

The content of each tweet can consist of text and additional, interactive content. Table 2 shows statistics on the usage of different content types.

The most prominent type is *user mentions* (85M). Since every retweet, reply, and quote contains at least one mention to the originator, these automated user mentions make up for 35%, 29%, and 3%, respectively. Therefore, 33% of the 85M mention-objects (28M) are user mentions, which are added manually into a tweet (*@username*). *URLs* (18M) are the second most prominent objects found in 23% of all tweets. There were 6 667 962 distinct *URLs* shared that originated from 275 078 different domains. Since 1/4 of all users in our corpus actively shared at least one URL, it seems typical for the German user base to consume and share content from external sources.

Beside these external sources we extracted 19.7 million (5 874 013 distinct) multimedia-objects. The majority of the multimedia contents shared are photos (82%), followed by videos (12%) and animated GIFs (4%), shared by a total of 56% of the users. Note that we can only obtain multimedia content from text tweets, as at least a single word is needed to identify a tweet to be German. Further, 29% of the users in our data set shared 39 million hashtags in 22% of all tweets. However,

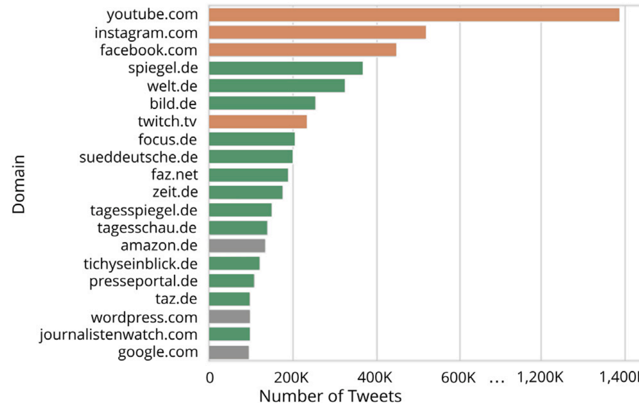

Fig. 3. Tweet volume of top 20 external sources (orange: OSN, green: news content, gray: other).

while we observe more tweets with hashtags than multimedia content, more users share multimedia content (56%) than hashtags (22%). Users using hashtags are about two times more active on Twitter than users sharing multimedia content, which, to some extent, explains this effect. A feature almost entirely neglected by users in our data set is the submission of geolocation data (*Places*). Only 2% of the users share their location when tweeting.

We want to understand the type of content circulating in the German-speaking Twitter community and measure the share of news-related contributions. Leveraging hashtags, shared external content, and the concept of FGs, we report on the media-consuming behavior of the German Twitter population.

Hashtags. In addition to external sources, users produce a high amount of hashtags. By examining popular hashtags shared during unusual high peaks in daily usage, we could identify the related influential events (see Table 4). We observe a peak in activity at the end of the 2019 European Parliament election. The election and discussion on the results dominate the hashtags during that time.

We also observe that hashtag usage does not reflect election results. The far-right party *Alternative für Deutschland (AfD)*, for instance, is close to leading the hashtag ranking (#AfD), even though it only came in 4th place in the election. Besides political events, pop-cultural events also caused an increase in daily Twitter volume, e.g., a non-German hashtag referencing the Korean pop band *BTS* or *Germany's Next Topmodel* (#GNTM) and the *Eurovision Song Contest* (#ESC2019). Here, the band *BTS* achieved high music chart rankings over several weeks in Germany, released a single, and, thereby, generated several trending hashtags. Nevertheless, most top hashtags correspond to events within German-speaking countries. These events also dominated the news in Germany during the data collection period.

Functional groups. Table 5 details the distribution volumes of the Top 10 FGs and their categories. We discuss the statistics in detail in Appendix A. Here, we focus on news related FGs. The most traffic is generated in the *Information/Communica-tion* FG (47% tweets). Large amounts of its content is related to news (General News: 32%) and personal blogs (Blogs/Wiki: 10%), mainly consisting of content from online news media and personalized political websites. Based on the high number of retweets in this group (52%), news and blog content seems to be well-received by the user base. We observed the same popularity of political domains in the FG *Society/Education/Religion*, comprised of even more elaborate political content. McAfee's TrustedSource grouped *Controversial Opinions* into the FG *Lifestyle*. The number of retweets in this category is 70%, further supporting the assumption that political content on Twitter is widely distributed and acknowledged.

Table 5
FG and categorical distribution of the 17 million URL-Tweets (multiple assignments per domain possible) and form of distribution, such as Original Tweets (OT), Retweets (RT), Replies (RP), and Quotes (QT); statistics on third-party services (Third) are included; FGs and categories with less than a 1% share of all tweets are excluded; FGs containing categories of the ***News Group*** are depicted in brown.

Category	Tweets %	Users %	URLs %	OT %	RT %	RP %	QT %	Third %
Information/Communication	**47**	**35**	**39**	**46**	**52**	**2**	**1**	**29**
General News	32	21	23	40	57	2	1	23
Blogs/Wiki	10	16	10	52	44	3	1	36
Public Information	2	3	2	68	29	3	1	59
Portal Sites	2	5	2	46	52	2	1	20
Technical/Business Forums	1	2	1	66	31	2	0	52
Forum/Bulletin Boards	1	2	1	64	33	3	0	43
Entertainment/Culture	**15**	**43**	**13**	**44**	**50**	**5**	**1**	**21**
Streaming Media	10	36	8	42	52	6	1	17
Media Sharing	8	33	6	39	54	7	1	14
Entertainment	4	10	4	56	41	2	1	38
Internet Radio/TV	1	1	0	69	29	2	1	53
Art/Culture/Heritage	1	2	0	37	60	2	1	21
Lifestyle	**12**	**23**	**17**	**65**	**34**	**1**	**0**	**55**
Social Networking	7	18	12	69	30	1	0	64
Sports	3	4	3	66	32	1	1	49
Controversial Opinions	1	1	0	29	70	1	0	11
Travel	1	1	1	84	13	3	0	73
Society/Education/Religion	**9**	**13**	**7**	**35**	**59**	**6**	**2**	**17**
Politics/Opinion	3	5	1	27	68	4	2	14
Education/Reference	2	5	2	43	46	9	3	23
Non-Profit/Advocacy/NGO	2	6	2	35	60	4	3	10
Government/Military	1	3	1	30	61	8	4	16
Health	1	1	1	52	41	5	2	35
Purchasing	**8**	**10**	**10**	**75**	**22**	**3**	**1**	**58**
Marketing/Merchandising	3	5	4	73	24	3	1	59
Online Shopping	3	4	3	71	25	3	0	52
Auctions/Classifieds	1	1	1	91	9	1	0	57
Business/Services	**6**	**10**	**8**	**71**	**26**	**2**	**1**	**52**
Business	4	8	5	65	31	3	2	42
Finance/Banking	1	2	2	78	20	2	1	63
Job Search	1	1	1	92	8	0	0	86
Information Technology	**5**	**12**	**7**	**69**	**28**	**3**	**1**	**56**
Internet Services	3	7	4	73	24	2	1	60
Software/Hardware	1	3	2	83	15	2	0	72
Pornography/Nudity	**4**	**5**	**3**	**43**	**56**	**1**	**0**	**43**
Pornography	2	2	2	49	51	0	0	63
Incidental Nudity	2	3	1	35	64	1	0	19
Games/Gambling	**3**	**5**	**2**	**57**	**42**	**1**	**1**	**38**
Games	3	5	2	57	42	1	1	37
Risk/Fraud/Crime	**1**	**1**	**1**	**65**	**33**	**2**	**0**	**77**

4.1.3. External media usage

A closer look at the 20 most shared external sources (see Fig. 3) revealed that 13 link to popular German news providers such as *Spiegel*, *Welt* or *Bild*, as well as to smaller news/opinion blogs, such as *Tichy's Einblick* and *Journalistenwatch*.

However, it turned out that the top domains are external OSNs, led by YouTube, followed by Instagram and Facebook (see Fig. 3). These platforms have a significantly higher distribution and more users sharing content from these platforms than any other domain (see Tables 6 and 7). They are platforms for a variety of content providers. Therefore, we resolved links to *YouTube*, *Facebook*, and *Instagram* to identify popular *YouTube Channels*, *Facebook Pages*, and *Instagram profiles*.

YouTube content varies from music, gaming, and political opinions to educational content (see Table 8). We identified single videos accounting for large chunks of the YouTube links on Twitter. For example, a newly released single of a Korean pop band (BTS) or a video of a channel called *Rezo* belonging to a person who was at the center of a political controversy surrounding the 2019 European Parliament election. He published a video with the title "Die Zerstörung der CDU" (Engl.: the destruction of the CDU) that went viral, expressing concern regarding the political course of the *CDU*. In general, there is only a small number of frequently shared content providers from YouTube (see Table 9). Half of these Channels are related to political topics. Moreover, they show a specific political affiliation. Channels belonging to the right-wing political party AfD are shared more often than channels of any other party. This observation indicates a high activity during their election campaign and shows a trend towards utilizing multimedia content to reach a broader spectrum of users.

Table 6
List of popular services used to distribute social media URLs.

Platform/App	Tweets	Users	Official
YouTube			
Android	33	48	✓
Web Client	26	17	✓
iPhone	17	27	✓
Web App	6	6	✓
IFTTT	3	0	✗
Instagram			
Instagram	62	42	✗
Android	14	25	✓
iPhone	8	17	✓
IFTTT	6	3	✗
Web Client	5	9	✓
Facebook			
Facebook	56	54	✗
Android	16	17	✓
Web Client	10	12	✓
iPhone	8	11	✓
Web App	4	4	✓

Table 7
Top shared external OSN content providers broken down by tweet type.

Platform	Count	Original	Retweet	Reply
YouTube	1 402 441 (374 414)	38	55	7
Instagram	520 466 (370 510)	72	27	1
Facebook	454 128 (292 316)	67	31	1

Table 8
YouTube URLs: Most shared video categories.

Category	Share (%)	User (%)	Video count
Music	20	34	47 018
News & Politics	19	18	18 484
Gaming	14	10	40 427
People & Blogs	13	20	29 610
Entertainment	12	21	21 982
Education	4	7	10 274
Science & Technology	4	8	8 332
Film & Animation	3	7	7 967
Nonprofits & Activism	2	4	3 868

Table 9
Top ext. social media profiles (Brown: Political Emphasis).

Provider	Tweets #	Users #	URLs #	Category	Description
YouTube					
첵코리아	284 980	282 903	1	music	South Korean singer
Rezo ja lol ey	50 277	30 802	29	political cont.	Rezo controversy
AfD Kompakt TV	14 323	3 413	99	political party	Political party: AfD
Rammstein Official	11 793	8 664	70	music	German band
ProDogRomania e.V.	10 265	617	637	activism	Dog rescue Romania
AfD-Fraktion Bundestag	8 201	2 351	309	political party	Political party: AfD
ibighit	8 094	7 328	36	music	Korean pop band
Joko & Klaas	7 540	6 231	35	entertainment	German entertainers
RT Deutsch	7 138	2 321	1 003	news/politics	Russian news media
Gottfried Curio	6 904	2 330	50	politician	Politician from AfD
Instagram					
@zkdlin	16 845	6 933	202	music	South Korean singer
@oohsehun	9 752	7 102	91	music	South Korean singer
@ksh7909	4 342	3 878	4	music	South Korean singer
@sooyoungchoi	3 579	1 720	9	music	South Korean singer
@daniel.k.here	3 093	3 038	9	music	South Korean singer
@taeyeon_ss	2 666	1 155	22	music	South Korean singer
@saulami1g	2 453	7	2 443	gaming	Gaming/Streaming
@svchicas	2 113	219	2	nudity	Explicit Content/Spam
@stephenathome	1 957	1 953	5	politics	Late night show host
Facebook					
@aliceweidel	16 433	3 867	327	politician	Party member of AfD
@alternativefuerde	11 762	2 930	190	political party	Facebook page of AfD
@Prof.Dr.Joerg.Meuthen	8 149	2 496	85	politician	Party member of AfD
@Bjoern.Hoecke.AfD	3 498	1 523	54	politician	Party member of AfD
@Pazderski.Georg	2 408	910	59	politician	Party member of AfD
@Academia-Para-C...	2 131	222	1	nudity	Explicit Content/Spam
@Deutschland3000	2 012	1 988	15	education	Educational (politics)
@GegenDieAfD	1 740	825	140	activism	Activism against AfD
@GottfriedCurio.AfD	1 694	1 138	10	politician	Party member (AfD)
@app: rossmann.de	1 472	1 079	1	advertisement	Facebook app (shop)

Instagram links are mostly apolitical and dominated by profiles from the entertainment industry. Looking at the most shared Facebook profiles (see Table 9), we observe a relatively small user base that only supports a handful of Facebook pages or profiles, with a low distribution factor. We notice, however, that most Facebook profiles are politically motivated and shifted towards the right-wing party AfD. One exception to this rule is a frequently shared page that directly opposes said party (@GegenDieAfD).

Overall, the top content providers from YouTube and Facebook are mostly related to political parties and activism. We observed that the German political party *AfD* was highly active on social media. Regarding shared links from Facebook (6 out of 10) and YouTube (3 out of 10), *AfD*-related topics dominated this content.

In general, with regard to the total number of tweets, only *BTS* and *Rezo* were able to generate reach comparable to other popular content providers on Twitter. For further discussions on the content from YouTube, Instagram, and Facebook see Appendix D.

Table 10
Hierarchical partitions of the Louvain Method of the unweighted (w^-) and weighted ($^+$) graph; CU depicts the number of communities with controversial users.

Level	Modularity Q		# of Communities	
	w^-	w^+	w^+	CU
0	0.398	0.7300	872 581	1 734
1	0.496	0.8800	334 656	435
2	0.529	0.9040	274 689	162
3	0.532	0.9056	270 437	117
4	0.532	0.9057	270 184	112
5	0.532	0.9057	270 177	112

Table 11
News Group Volume: Number of URL- and Reaction-tweets.

Data set	Tweets		Users		URLs	
	#	%	#	%	#	%
URL-tweets	17 478 261	100	1 720 752	100	6 667 962	100
News Group	7 247 843	41	454 381	26	1 903 133	29
Reaction-tweets	9 582 682	100	1 222 863	100	1 193 232	100
News Group	5 660 382	59	391 139	32	515 883	43

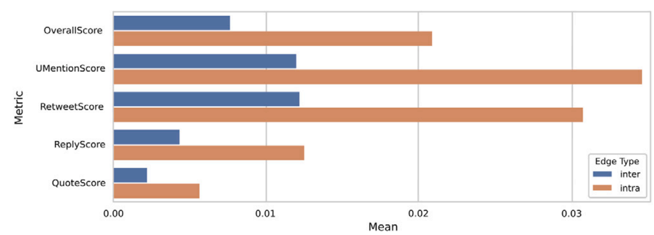

Fig. 4. Comparison of the interaction metrics based on edges within communities (intra) and edges between communities (inter).

4.1.4. Community structures

We further explore user behavior related to political discussions. Therefore, we study the community structures of the German-speaking Twitter community. Statistics on activity, tweeting behaviors, and communication allow us to analyze group dynamics and -characteristics.

Our studies are based on a holistic interaction graph that stems from 29 098 133 retweets, 24 432 025 replies, 2 907 173 quotes, and 37 979 345 mentions to users within the network. The final graph encompasses 6 809 903 users connected via 32 984 267 edges.

In the list of identified communities, we observe a consistent number of groups with 2–3 members (2: 206 054, 3: 38 551). Due to their inactivity, these tiny groups do not get merged into larger communities. We consider these as noise. They interact with one or two other users and do not contribute to conversations and controversies. The remaining community structure of our Twitter corpus encompasses 25 572 communities. In the following, we study the quality of detected communities according to the mentioned quality indicators (see Section 3.4.3).

Table 10 reports on the modularities of detected communities (weighted- and unweighted). The modularities are $Q = 0.53$ and $Q = 0.91$, respectively. Thereby, they exceed the expected maximum modularity for a random network ($Q = 0.15$), suggesting reliable community structures. The modularity score of 0.91 indicates that the discovered communities describe groups of users having a significantly higher exposure to each other than to users from other groups.

Table 10 summarizes the 5 iterations of the community detection approach. It depicts 6 partitions of communities with increasing modularity scores. In the first iteration, the unweighted graph barely reaches the acceptance threshold of the modularity score of 0.4. Acceptable modularity scores in the initial aggregations are advantageous. Including the lowest hierarchical level in the analysis circumvents the

Table 12
Statistics on identified promotional profiles (left) and reach (right) of the most shared content providers within the News Group (U = Users, J = Journalists); well-respected traditional German news providers (acc. to [55]) are highlighted in blue; for corresponding URLs see Table 18.

Content prov.	U	J	Feeds	Tweets	URLs	OT	RT	RP	QT	3rd	UTweets	Users	U/T	URLs	OT	RT	RP	QT	3rd
	#	#	#	#	%	%	%	%	%	%	#	#	#	#	%	%	%	%	%
Spiegel	98	68	30	20 850	24	86	13	0.42	0.18	72	344 946	76 028	4.54	47 805	31	66	2	1	12
Welt	72	53	19	23 017	39	88	12	0.23	0.08	90	302 259	47 139	6.41	43 842	26	71	3	0.40	7
Bild	138	59	79	31 059	52	97	3	0.05	0.01	94	221 394	30 547	7.25	24 526	21	78	1	0.21	5
Sueddeutsche	76	38	38	14 231	27	91	8	0.43	0.16	81	184 966	55 572	3.33	20 990	24	74	2	0.43	9
Zeit	80	59	21	7 724	24	88	11	1	0.27	82	163 648	51 052	3.21	21 314	23	73	4	1	9
FAZ	59	34	25	26 322	35	94	6	0.16	0.08	89	163 531	40 784	4.01	32 909	28	70	2	1	8
Focus	20	5	15	49 952	53	81	19	0	0	81	154 124	23 751	6.49	25 050	27	71	2	0.25	11
Tagesschau	10	8	2	2 463	21	98	2	0.04	0.04	95	138 522	39 178	3.54	10 770	27	71	2	0.46	14
Tagesspiegel	77	57	20	12 559	38	45	53	1	2	0	138 274	38 231	3.62	13 590	18	79	2	1	6
Tichys Einblick	2	1	1	1 483	29	88	9	1	1	15	120 412	9 344	12.89	2 315	8	91	1	0.34	3
Presseportal	3	1	2	5 550	10	100	0	0	0	100	103 418	10 401	9.94	52 888	55	44	1	0.48	43
Journalisten…	1	0	1	2 151	56	100	0	0	0	100	95 225	6 470	14.72	3 789	22	78	0	0.04	3
taz	43	26	17	6 754	44	91	8	1	0.07	83	91 674	29 327	3.13	8 432	16	82	2	1	7
Heise	26	14	12	4 946	23	76	24	0.22	0.14	93	90 829	25 667	3.54	15 330	38	58	3	2	22
n-tv	12	5	7	5 177	24	95	4	0.19	1	85	88 036	21 254	4.14	19 304	32	66	2	0.39	13
NZZ	82	55	27	10 395	35	58	38	4	0.14	20	78 256	22 146	3.53	14 883	33	64	2	0.39	11
Handelsblatt	66	60	6	13 090	36	74	25	0.34	0.11	60	76 174	22 765	3.35	23 212	38	60	2	0.44	19
Epochtimes	2	1	1	192	2	100	0	0	0	100	73 086	6 898	10.60	8 313	25	74	1	0.13	6
Philosophia p…	0	0	0	0	0	0	0	0	0	0	72 926	6 408	11.38	1 457	19	81	1	0.23	2
ZDF	64	35	29	3 877	29	80	14	5	0.48	19	70 492	27 887	2.53	9 914	18	79	2	1	6
Der Standard	33	28	6	8 131	39	96	4	0.33	0.09	89	66 846	15 745	4.25	14 495	37	60	3	1	13
change.org	11	8	4	169	0	51	31	14	4	3	61 586	27 915	2.21	34 674	59	39	3	1	3
Junge Freiheit…	2	1	1	672	26	83	17	0	0	1	56 823	7 026	8.09	1 864	14	85	1	0.07	5
Deutschlandf…	4	3	2	3 824	31	86	9	5	0.18	4	49 799	19 815	2.51	9 178	25	71	3	1	9
WDR	38	24	16	3 966	36	72	15	9	5	15	44 134	18 345	2.41	5 218	18	80	2	1	7
bundestag.de	4	1	3	71	1	100	0	0	0	8	43 280	19 733	2.19	4 233	17	76	6	6	9
BR	59	40	20	8 937	63	87	10	2	1	47	42 108	16 190	2.60	7 805	23	74	3	1	9
Stern	17	10	7	5 588	29	98	2	0.18	0.07	93	41 611	14 892	2.79	16 231	43	55	2	0.46	27
NDR	37	22	15	3 860	32	64	34	2	1	9	40 851	16 025	2.55	6 851	28	70	2	1	13
RT	10	0	10	2 334	40	91	1	7	0	0	39 262	6 683	5.87	5 189	31	67	2	0.31	7

resolution limit of modularity optimization [47]. Lancichinetti and Fortunato confirmed this in a comparative analysis of community detection methods, using the lowest hierarchical level to improve their performance [54]. Therefore, during the analysis, we take every partition into account. Fig. 4 depicts the expected values of the different interaction metrics (final iteration). Scores of partitions from lower hierarchical levels show similar results but produce inter and intra-edges with slightly higher expected values.

These results suggest that the final partition represents a more generalized overview of the user groups. Partitions in earlier iterations, however, depict smaller communities in more detail. They allow for a better understanding of these groups within the network.

Finally, we consider the size distribution of the identified groups. We observe the desired skewed distribution in our community structure. 45% of our users belong to the 10 largest communities. The lowest level of the hierarchical partitions shows a flatter distribution with only 13% of the users belonging to the 10 largest communities.

4.2. News content analysis

We aim to investigate informational and political content on Twitter and how it influences the German user base. We defined the News Group as a collective term that comprises external domains related to news, political/controversial opinions, and educational content (see Section 3.3.1). In the following, we leverage our knowledge on shared content to further our understanding of news-related information.

Table 11 shows the volume of tweets, users, and URLs within the News Group. Approximately 41% of all URL-tweets distribute content that belongs to this group. However, only 26% of the users sharing URLs belong to this group. The ratio between URL-tweets (41%) and distinct URLs (29%) in the News Group implies that the average URL is shared 3.81 times. Compared to the average distribution of non-members with a distribution factor of 2.15, the News Group is more active in sharing the content of interest. Therefore, URLs shared on Twitter predominantly link news-related content.

4.2.1. News exposure

We established that 25% of the user in our Twitter corpus shared at least one URL-tweet, and 18% of the users replied. With 26% news-related URLs, we have 6.5% of the users actively sharing news content. However, these numbers are only a lower bound on the percentage of users exposed to external content. The challenge is identifying users that read and consume but do not react to URL-tweets. These users only use Twitter as a news-feed and show no measurable activity towards URLs at all. Although we cannot accurately estimate the number of these users, it is possible to identify their position in the network. The attempt we follow is to find the communities that share URL-Tweets. Since communities are densely connected, their users are also more exposed to content from within the community. Therefore, we assume that URL-sharing communities expose their members to external news sources.

Hence, by counting the members of (news-related) URL-sharing communities, we obtain an upper bound on users exposed to external content.

Overall, 23% of the communities share URLs. They encompass 91% of all users. 28% of these, 1678 communities, share news-related content. They still combine 90% of all users and produce 99% of all tweets. A mean percentage of 62% news-related links indicates that most URLs within these communities relate to news, political, or educational topics. With an average of 72%, the ratio within Reaction-tweets is even higher. In total, 35% of their users produced 34% of the tweets while immersed in news-related discussions. Projected onto the entire data set, 31% of all tweets in our data sample discuss external news content. These results suggest that URLs from external news sources attract more attention than any other content. Informational content has a massive influence on the German-speaking Twitter community.

4.2.2. Engagement

We established that 13 of the 20 most shared external sources (see Fig. 3) link to popular German news providers such as *Spiegel*, *Welt* or *Bild*, as well as to smaller news/opinion blogs, such as *Tichy's Einblick*

Table 13
Categorical usage and distribution of 7M URL-tweets (450k users), 6M reaction-tweets (390k users) within the News Group (URL/reaction); categories are distinguished by political views from: moderate- to extreme.

Category	Tweets %	Users %	Distribution (%) OT	RT	RP	QT	Third
General News	77 / 82	79 / 85	41	57 / 41	2 / 49	1 / 22	23 / 3
Politics/Opinion	8 / 8	19 / 20	27	68 / 47	4 / 44	2 / 25	14 / 3
Education/Reference	5 / 4	20 / 15	43	46 / 36	9 / 54	3 / 22	23 / 4
Non-Profit/Adv./NGO	5 / 3	21 / 14	35	60 / 48	4 / 40	3 / 32	10 / 4
Controversial Opinions	3 / 3	2 / 3	29	70 / 68	1 / 25	0 / 14	11 / 1
Government/Military	3 / 3	13 / 14	30	61 / 43	8 / 46	4 / 34	16 / 3
Major Global Religions	1 / 1	3 / 3	42	54 / 35	4 / 54	2 / 20	20 / 3
Discrimination	<1 / <1	<1 / <1	42	32 / 51	24 / 41	2 / 12	5 / 1
Historical Revisionism	<1 / <1	<1 / <1	59	19 / 60	22 / 32	<1 / 34	2 / <1

and *Journalistenwatch*. To further our understanding of news distribution within the German-speaking Twitter community, we analyze the subset of shared external sources that link to news-related content. Subsequent analyses are based on the 30 dominant news providers in our data set.

We turn our attention to user engagement. We compare the popularity and reach of content providers within the News Group by analyzing the volume of tweets they generated, the number of users they mobilized, and the number of reactions they prompted.

With the bouquet of actors and news providers, we shed light on the most influential distributors and how users support and react to these diverse options. Regarding controversial content, we further analyze the influence on the general public (on Twitter).

We study user engagement by measuring two factors: reach and impact. We approximate reach by the spread of URLs from a content provider and calculate impact by the number of reactions towards these tweets. In the following, we report on reach and impact w.r.t. two different aspects: (*i*) category, and (*ii*) news provider. We also cover engagement towards links of external OSNs.

Reach. The *News Group* comprises 9 categories. These categories allow us to examine the reach w.r.t. different types of news. Table 13 gives an overview of the sharing behavior viewed by category. Most URL-tweets originate from moderate domains (General News: 77%). Besides religious- (20%) and educational content (23%), general news is with 23% on the top of the list w.r.t. the distribution via third-party services. Regarding support via retweets, we observe that news sources that offer tendentious to extreme views on politics (i.e., Politics/Opinion and Controversial Opinion) are the most supported domains (retweet factor: 68%–70%). However, the average distribution of URLs via retweets is consistently high in almost all categories. An exception is URLs propagating extreme political views, i.e., discrimination and historical revisionism. With a retweet factor of 19%–32%, such content experiences significantly less support via retweets. Interestingly, however, these links seem to be often used within discussions, resulting in a 22%–24% URL-tweet share via replies (others: 1%–9%).

The data suggests three different support patterns: (*i*) highly shared and discussed articles, (*ii*) highly distributed articles via retweets (68%–70%), and (*iii*) articles supported via replies (22%–24%) but mainly ignored by the general public ($\leq 3\%$ of all users). These support patterns correlate strongly with the subjectivity level of shared content, i.e., *moderate* domains are supported by (*i*), *tendentious* outlets by (*ii*), and *extreme* domains by (*iii*). Overall, we rarely observe extreme external content. A share of $< 4\%$ of URL-tweets, actively shared by $< 4\%$ of the users, and rarely replied to, extreme content seems to be widely ignored by most Twitter users.

Concentrating on news providers, Fig. 5 depicts the tweet volume broken down by provider. Further, Table 12 (right column) shows additional data regarding tweet distributions. The user/tweet ratio reveals two distinct user types. Followers of domains such as *tichyseinblick.de*, *journalistenwatch.de*, or *philosophia-perennis.com* (tendentious to extreme views) have the most active users with a user to tweet ratio of 10.83 (tendentious) and 12.20 (extreme). In comparison, readers

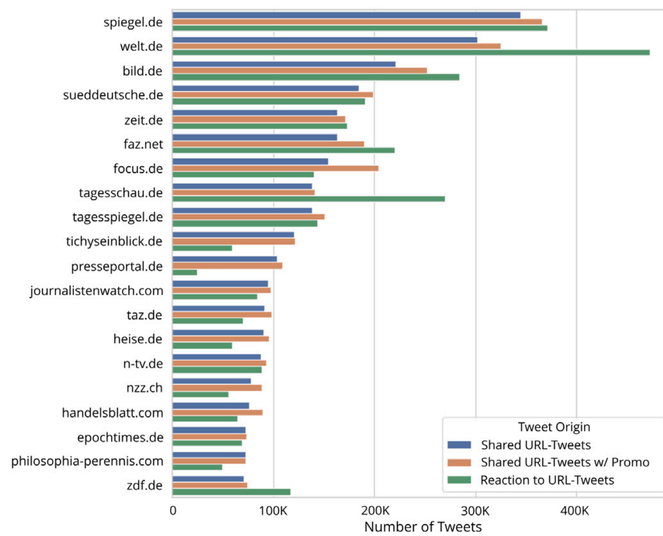

Fig. 5. Tweet volume of the top 20 news domains, comparing URL-tweets (with and without promotion profiles) and Reaction-tweets.

of traditional news outlets such as *Spiegel* or *Zeit* only have a ratio of 4.54 and 3.21, respectively (traditional German news providers: 3.91). Also noticeable, Twitter users sharing less moderate outlets retweet more often, with *tichyseinblick.de*, *jungefreiheit.de*, *taz.de*, and *philosophia-perennis.com* as top domains in this category and retweet counts ranging from 81% to 91%. Note that similar to *Bild*, while *taz* is part of moderate news media, in the past several articles with tendentious, disputable content were rebuked by the German Press Council.[4] Traditional media sources, in contrast, reach a broader spectrum of users, but their popularity partially depends on the number of articles they publish.

Users neither share the links from moderate nor tendentious media via replies, indicating that users less often reference such content within discussions. In this category, the governmental outlet *bundestag.de* shows the highest reach in this context with distributions via replies and quotes of 6%, each.

We complement information on the reach of news providers by studying promotional profiles, i.e., we consider self-promotional tweets produced by feed-profiles and corresponding journalists. Table 12 (left column) details the results of our promotional profile detection process for each of the 30 content providers. For instance, we identified 98 promotional profiles from *Spiegel*, comprised of 68 journalist-profiles and 30 feed-profiles. Over two months, these profiles produced 20 850 tweets, which constitutes a daily average tweet volume of ~ 342 tweets (per account: ~ 3.49). These accounts mainly distributed content via

[4] https://de.wikipedia.org/wiki/Die_Tageszeitung#Presseratsr%C3%BCgen

Table 14
Reaction-tweets towards the 30 most distributed content providers from the News Group.

Provider	Tweets		Users		URLs		Distribution (%)			
	#	%	#	%	#	%	RT	RP	QT	Third
Spiegel	371 725	10	72 881	14	17 884	35	36	56	20	2
Welt	472 260	11	62 868	15	24 592	46	37	56	19	2
Bild	283 968	10	43 158	13	14 790	48	40	52	15	1
Sueddeutsche	1 90727	9	53 591	12	9 088	40	37	53	25	3
Zeit	173 291	9	45 878	12	9 024	41	31	60	21	2
FAZ	220 151	10	46 895	13	14 820	40	34	57	23	2
Focus	140 549	7	23 089	15	8 960	19	46	46	15	1
Tagesschau	270 323	7	50 758	10	4 623	43	41	53	18	2
Tagesspiegel	144 017	8	37 048	12	7 176	50	37	53	27	2
Tichys Einblick	59 531	3	10 331	9	983	42	42	48	20	1
Presseportal	24 369	5	9 297	9	5 233	10	41	49	24	5
Journalistenwatch	84 145	8	7 616	10	2 450	64	68	25	14	1
taz	69 589	8	24 225	9	4 547	50	38	51	32	3
Heise	59 171	10	18 783	12	5 024	31	63	30	15	11
n-tv	88 275	11	23 658	13	8 027	39	44	48	16	2
NZZ	56 114	10	18 572	13	5 306	31	39	51	23	2
Handelsblatt	64 816	10	22 968	13	7 456	26	38	50	27	2
Epochtimes	68 923	7	7 542	14	3 144	37	63	30	18	0
Philosophia perennis	49 726	6	7 558	11	721	49	68	26	12	1
ZDF	117 548	8	31 993	8	4 559	44	33	58	23	2
Der Standard	60 219	13	14 356	14	6 770	37	47	44	23	3
change.org	17 808	6	9 697	8	3 277	9	55	38	20	3
Junge Freiheit	45 106	5	8 928	10	752	40	41	50	21	1
Deutschlandfunk	49 611	11	18 625	11	4 290	45	31	57	24	3
WDR	40 341	9	16 336	9	2 525	45	38	50	30	3
bundestag.de	26 863	6	12 334	8	1 135	27	45	44	40	3
BR	43 720	11	16 508	10	4 325	38	44	46	28	3
Stern	46 682	13	16 232	9	5 411	29	32	60	15	2
NDR	36 454	10	15 287	11	2 948	41	40	50	22	3
RT	36 089	12	7 524	11	3 193	58	58	34	18	1

original tweets (86%) and shared them via third-party services (72%). In the process, they actively distributed 27% of the distinct URLs from *spiegel.de* shared during the two months.

In general, we observed that predominantly traditional news media sources, such as *Spiegel*, *Welt*, *Bild* and *FAZ* disseminate their articles via third-party services to extend their reach on Twitter. In particular, *Focus* utilizes a sophisticated feed-profile network that produces a massive volume of tweets. Besides *Bild* with 52%, *Focus* also covers (53%) most of their articles circulating on Twitter, only topped by *Journalistenwatch* (56%) and *BR* (64%). In contrast, non-commercial public news media such as *Tagesschau* and governmental news providers such as *bundestag.de* only generate small amounts of such tweets utilizing significantly more diminutive Feed-networks.

We observed that tendentious to extreme outlets, such as *Tichys Einblick*, *Philosophia perennis*, *Journalistenwatch* and *Epochtimes*, generate much less self-promotional tweets than traditional media. Note, however, that these findings could also be an artifact due to our detection approach, i.e., the corresponding promotional profiles could not adhere to media best practices (see Section 2.3).

Impact. Besides reach, we measure the impact of news categories, -providers, and external OSNs. Table 11 shows the number of tweets commenting on or referencing URL-tweets. We found that most reaction-tweets (59%) occurred in the News Group. Furthermore, 43% of the URLs that prompted reactions on Twitter originated from the News Group. The proportion of users (32%) and tweets (59%) indicates a highly active News Group.

We analyzed the distribution of reaction-tweets considering each category of the News Group (see Table 13). In contrast to the distribution of URL-tweets, we registered almost no Reaction-tweets from third-party services. Regarding discussions, users commented on moderate content actively (replies + quotes: 71%–80%), followed by tendentious articles (69%). The more extreme the content, the more "discussions" via retweets (extreme content: > 50%) with an active discussion ratio (replies + quotes) of 53%–66%. These results suggest that some users continue to disseminate and support controversial opinions, while others are less likely to respond to such content (reply rate: General News 49% vs. Controversial Opinions 25%).

In terms of activity levels, it can again be seen that users discussing extreme content are the most active, with a ratio of tweets per user of 8.93. Users discussing tendentious content are this time more similar to users discussing moderate content, with 5.43 and 4.84 respectively. We find that users are more active in discussing than sharing moderate content 4.84 versus 3.91. The reverse is true for tendentious (5.43 vs. 10.87) and extreme (8.93 vs. 12.20) content.

Throughout, we observe many replies and quotes. Therefore, we assume that users heavily engage in political discussions. Note that the cumulative percentages of retweets, quotes, and replies exceed 100% because retweets may contain nested quotes and replies, which we counted as a retweet of each instance in this case.

Fig. 5 depicts the URL-tweet volumes and the number of Reaction-tweets concerning each content provider. Traditional news media, such as *Spiegel*, *Welt*, and *FAZ*, trigger a high amount of Reaction-tweets, exceeding the number of tweets that share their articles. It suggests that users actively discuss their content. Of note are the statistics of *Tagesschau*. While showing only moderate amounts of URL-tweet shares, it prompted the 4th-highest number of Reaction-tweets. Not reaching the number of their respective URL-tweets, tendentious and extreme media providers, in contrast, receive fewer reactions.

Table 14 gives a more detailed overview of the reactions prompted by the 30 most shared domains in the News Group. For instance, *Spiegel* received 371 725 Reaction-tweets to 10% of tweets that shared a *Spiegel* article. In total, 72 881 users reacted to 17 884 distinct URLs from *Spiegel*, encompassing 35% of all unique *Spiegel* URLs shared on Twitter. Thereby, only 14% of the users that shared a *Spiegel* article received any reaction.

The outlets of *Welt*, *Bild*, *ZDF*, and *Tagesschau* trigger a significantly larger user base discussing their content than the user base that shares it.

Finally, we explore content from external OSNs. We start by further measuring the reach of content by the ratio of shares per unique URL

(S/U). URLs of the 10 most shared content providers reach an average S/U ratio of 12.58. YouTube- (S/U: 3.75), Instagram- (S/U: 1.40) and Facebook links (S/U: 1.55) were shared less often, and, hence, failed to generate reach and impact. A portion of 55% retweets and 7% replies when sharing YouTube-URLs suggests that users distribute YouTube videos to support the content and communicate with other users. On the other hand, the Twitter user base widely ignores Facebook and Instagram URLs. These links get mostly shared via original tweets (Instagram: 72%, Facebook: 67%). While users distribute YouTube links primarily via the official mobile and web clients from Twitter, they share most Facebook and Instagram content via third-party services. We assume that most Facebook and Instagram users share their content passively while actively using Facebook and Instagram clients. They share content on these platforms and forward them to their Twitter profiles to extend their reach. However, the low number of retweets (Instagram: 27%, Facebook: 31%) indicates that this strategy is not very effective. Consequently, we conclude that Facebook and Instagram content is perceived less distinctly than YouTube or other shared media content. Overall, compared to news media sources that distribute their articles directly on Twitter, content providers that operate from other social media networks attract considerably less attention.

So far, our analyses on shared content (Section 4.2) and the engagement within the News Group (Section 4.2.2) showed that political content produces the most activity within the German user base. Further, the high number of reaction-tweets to URL-tweets from the News Group suggests keen interest in such content and that users use Twitter as a platform for political discourse. Regarding political content, non-controversial content attracts more users than controversial topics. Traditional news providers distribute most of the news articles (see Fig. 3). However, articles from tendentious to extreme outlets generated significantly more retweets per user. Still, the majority of users support moderate views (via retweets). Only a small group ($<$ 4% of the users) supports extreme political views. These users tend to use the reply functionality instead of retweets, implying that they share their opinions within discussions.

4.2.3. Political hashtags

Next, we analyze hashtag-usage w.r.t. categories to further our understanding. Table 15 gives an overview of popular hashtags within news categories and compares URL- and Reaction-tweets. A variety of content providers report on the same events. Therefore, many popular hashtags, such as #AfD, #Europawahl2019, and #Rezo, appear in multiple categories. We also observe that most of the hashtags in General News exhibit a political background. Based on the reaction-tweets prompted by these categories, we observe that users often reply to political news concerning the CDU with hashtags that dissent the party and its coalition partner (e.g., #NieMehrCDU, #NieMehrSPD). Reaction-tweets indicate that users discuss the shared news and use hashtags to express their opinion. While many news articles express less extreme opinions, reaction-tweets express their views more directly. Tweets that distribute controversial news content also receive attention from users with opposing opinions, observable by the usage of hashtags like #MeinungsfreiheitAuchFürDumme (Engl: free speech even for idiots) and #homophob within the reaction-tweets. Users sharing discrimination sources use hashtags opposing the *CDU* and *SPD*. In contrast to other categories, Reaction-tweets contain fewer opposing hashtags. Only a minor fraction of the user base discusses content from discrimination sources without attracting much attention from users opposing their views.

4.2.4. Communities

We complement our studies, including information on community structures. Table 17 provides detailed information on the 20 largest communities. With 62% of all tweets in our Twitter corpus, the largest community substantially determines the content we observe in the German-speaking Twitter community. We reference this community, comprised of 816 677 users, as the German Twitter Core Community (Core).

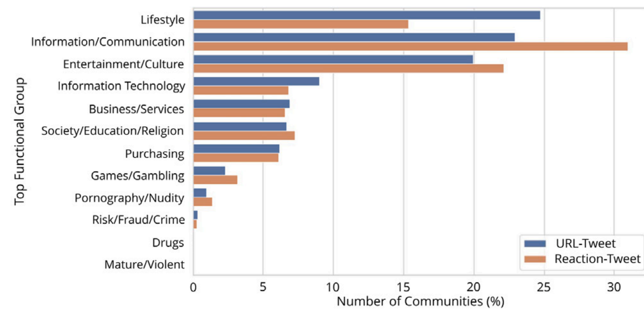

Fig. 6. Most shared Functional Groups.

Further, since we identified many communities, it is sensible to obtain an overview of popular users and hashtags. Although not detailed enough to understand the content discussed within a cluster, it provides a high-level approximation. We report on the 10 largest communities, including the most popular users and hashtags (see Table 16). User popularity is measured by PageRank, indicating how well-connected someone is.

Popular hashtags. Table 16 shows that most of the Core's top hashtags are related to politics (e.g., #AfD, #Europawahl, etc.) and activism (e.g., #FridaysForFuture). Moreover, users like Rezo (@rezomusik) and A. Kramp-Karrenbauer (@akk) also relate to political events. On the other hand, we found multiple clusters engaged in Asian pop-culture personalities and events. We examined the community 109 805 and found numerous users distributing music and entertainment content. By looking at the top users and hashtags, we can assess the general orientation of many communities. For example, the community 262 453 shows several gaming-related profiles and hashtags, whereas 150 111 and 142 499 are related to a mix of lifestyle topics and hobbies. With community 257 645, we discovered a political group exclusively discussing non-German content. The central users within this community are related to US politics, including Donald Trump (@realDonaldTrump) and his daughter (@IvankaTrump, @FLOTUS). We also found that mainly international media sources and YouTube videos are distributed in this community. International communities do not necessarily mean that they contain no German users but that their interests include global content.

Popular FGs. By measuring the most popular FGs, we identify the dominant domain category of a community (see Fig. 6). We report on 5860 communities containing users sharing URL-tweets. We can observe that most communities evolve around external sources classified as Lifestyle, Information/Communication, and Entertainment/Culture. Spam-like external sources, such as Pornography/Nudity, Risk/Fraud/Crime, and Drugs, only dominate a fraction of communities. The relatively low percentage of communities that react to Lifestyle sources confirms our assumption that users mainly ignore these tweets, including links from other social media networks. In contrast, many communities predominantly react to Information/Communication sources, including news and blogs. Our approach of classifying communities based on their URL-sharing behavior is also sensible for filtering communities. For example, we detected numerous communities that only share spam URLs and inappropriate or malicious content.

Overall, communities significantly differ in tweeting behavior, interest, and connectedness.

Popular news categories. To extend our approximations on the popularity of news-related topics, we calculated their spread within the community structure. By observing the most shared domain category, we observed that most News Group communities (67%) engage in content from General News sources such as Spiegel, Welt,

Table 15
Popular Hashtags within the News Group.

Category	Hashtags (URL-Tweets)	Hashtags (Reaction-Tweets)
General News	AfD, SPD, Berlin, CDU, ots, news, Europawahl2020, Merkel, FridaysForFuture, EU, Europawahl, Deutschland, NotreDame, Klimaschutz, Polizei	AfD, SPD, CDU, Europawahl2019, Merkel, FridaysForFuture, EU, Berlin, NieMehrCDU, Deutschland, Europawahl, Rezo, Strache, FPÖ, Klimaschutz
Politics, Opinion	AfD, Europawahl2019, EU, Europawahl, PIRATEN, SPD, Europa, Bundestag, EP2019, CDU, Prüffall, Deutschland, FridaysForFuture, Liebe, ReconquistaInternet	AfD, Europawahl2019, CDU, SPD, NieMehrCDU, PIRATEN, Europawahl, TERREG, EU, Piraten, NieMehrSPD, Uploadfilter, FridaysForFuture, FDP, CSU
Education, Reference	FridaysForFuture, Rezo, Europawahl, Europawahl2019, Klimaschutz, FFFfordert, actnow, Digitalisierung, OSTSTEINBBEKKER音, GrimmsWort, OTD, Berlin, DOYOUNG, KI, Stellenangebot	FFFfordert, actnow, wespoke, OER, GoBlue, kangdaniel, 김종현, WelcomeBackDaniel, 임영민, ABSOLUTE6IX, Marburg, noplanetB, twitterlehrerzimmer, wählengehen, Twitterlehrerzimmer
Non-Profit, Advocacy, NGO	Europawahl2019, Rezo, Zensur, Homöopathie, Meinungsfreiheit, Klimaschutz, Europawahl, Europa, FridaysForFuture, Klimakrise, EU, Uploadfilter, AfD, Berlin, Transsexuellengesetz	Lifeline, Scientists4Future, Florida, unteilbar, Atheisten, AlleGegenRWE, Weimar, OperationSophia, SafePassage, GrandTheftEurope, Economists4Future, Garzweiler, Thema, Upskirting, GamerGate
Controversial Opinions	FFD365, AfD, anonymous, anonymousnews, NotreDame, Merkel, EU, SPD, Antifa, EU19, Berlin, Grüne, CDU, Papst, Migration	anonymous, OliverFlesch, RRG, anonymousnews, MiloYiannopoulos, ramadan, Sperre, Obdachloser, MeinungsfreiheitAuchFürDumme, Schönleinstraße, FFD365, Grosz, einschönesOsterfest, pädophil, homophob
Government, Military	Bundestag, AfD, keinluxus, Klimaschutzgesetz, Feuerwehr, Polizei, Klimaschutz, Europawahl2019, FridaysForFuture, ParentsForFuture, Petition, Fahndung, EU, Braunkohle, Urheberechtsreform	Urheberechtsreform, Feuerwehr, Urheberrechtsreform, BVerfG, Protokollerklärung, Fahndung, KeinAber, copyright, 1919LIVE, SPC_Watch, Vermisstenfahndung, txwx, NRWE, Barcelona, Rossell
Major Global Religions	Kirche, AfD, NotreDame, Frauen, Europawahl, Missbrauch, ZdK, Sternberg, Karwoche, Ostern, PapstFranziskus, Woelki, Europa, GehtWählen, Papst	Karwoche, BenediktXVI, Glaube, Ratzinger, Benedikt, Maria20, kirche, Tagesevangelium, klerikal, Kirchenkrise, Kirchenaustritt, Sexualität, berührende_Erzählung, Gründonnerstag, Freitagsworte
Discrimination	ISIS, falseflag, Churchill, H8Front, H84U, Weltkrieg, PeterPadfield, KJM, RudolfHess, niemehrCDU, niemehrSPD, sydney, kalergiplan, Gunskirchen, IMMIVASION	falseflag, ISIS, Afd, leftwing, Gruene, Gewalt, H8Front, H84U, Ibizagate, Linke, Podcast
Historical Revisionism	Churchill, Weltkrieg, Grundgesetz, PeterPadfield, RudolfHess, GG70, Freimaurerei, Verfassungsschutz, Kommunismus, 1Mai, Kühnert, Verfassung, Nationalsozialismus, Sozialismus	Grundgesetz, GG70, Euro, Verfassung, Verfassungsschutz, Verfassungsrichter

or FAZ. Additionally, we identified many communities sharing Education/Reference content (17%). These are relatively small in terms of user size. They frequently used sources regarding environmental activism and university-related information sites. Communities sharing Non-Profit/Advocacy/NGO sources (7%) show more ties to political activism and local charitable projects. Communities preferring Political/Opinions content (4%) support opposing views, e.g., sharing left-wing- (e.g., avaaz.org and campact.de) or right-wing sites (e.g., infowars.com). User discussing governmental/military content and groups discussing religious topics make up 4% and 1%, respectively.

News providers. Next, we refine our understanding of news providers by measuring their influence within the communities. Table 18 provides an overview of the tweet distribution within the network for each content provider. We distinguish between URL- and Reaction-tweet statistics. For example, *spiegel.de* is shared within 413 communities and reacted upon within 318. The table also provides the mean PageRank of the user base that shared the tweets. In the case of *spiegel.de*, the mean PageRank of its supporting users is in the 74th percentile. This percentile of the PageRank indicates that the average Spiegel reader is better connected than 74% of the users in our Twitter corpus. Furthermore, the mean degree shows the number of connections Spiegel readers have with other users in the graph. An average user who shares *spiegel.de* articles interacted with 261 other Twitter profiles during the two months of our data collection. We also observe that the users who reacted to *spiegel.de* articles are better connected in the graph (Mean PageRank: 78th percentile; Mean degree: 289) than people who share the articles. We observed the same pattern for most news media and political blogs. This finding indicates that people, who comment on news articles, are overly active on Twitter in general and better connected than users who only share URLs. Interestingly, readers of controversial media, such as journalistenwatch.com, epochtimes.com, and philosophia-perennis.com, are noticeably well connected on average.

Table 16
Popular users and Hashtags for the largest communities in our network. We also assigned different subjects to the communities based on our investigation of the respective data.

ID	Top User (PageRank)	Top Hashtags	Subjects
Core	*rezomusik, CDU, akk, janboehm, ChangeGER, SPIEGELONLINE, welt, tagesschau, faznet, DB_Bahn*	AfD, Europawahl2019, CDU, EU, SPD, FridaysForFuture, Berlin, Europawahl, Rezo, NotreDame	German, Politics, News, Election
109805	*BTS_twt, kbs_exclusive, nochuucometru, cookiesketches, jasonaron, H0NEYMYG, mygwithluv, uchiha_jungkook, AshToTheBashh, Taeholic_V*	방탄소년단, BTS, 뷔, 태형, BBMAsTopSocial, V, BTSV, taehyung, 방탄소년단뷔, 태태	Pop-Culture, Entertainment, Korea, Music
262453	*KPrime86, nusr_ett, Doodlelot, ikuchan_kaoru, pegushi_, _bazztek, Kitsune__Zakuro, modernmodels, steelix666_, Crystal_herb*	魔道祖, FFXV, FFBE, 天官福, nsfw, FFXIV, MoDaoZuShi, FGO, LeagueOfLegends, xenoblade2	Entertainment, Gaming
34263	*VXyeontan_, WayV_official, ShunJou, donlaima, petiteyoona, babykihyunie, itfeelquotes, __CONY13, The_LordOfSalem, OfficialMonstaX*	WayV, 威神V, WeiShenV, WINWIN, 董思成, 윈윈, TEN, 李永, ウィンウィン, NCT127	Pop-Culture Asia
249774	*yalibragirl, kellieeastwood, rihanna, __hazelr, ltsesthetics, femmeduart, itsbaddies, Stripx777, onlybaddies, RomeTrumain*	MetGala, WayV, 威神V, WeiShenV, TEN, 李永, The1975, WINWIN, 董思成, 뷔	International, Art, Entertainment, Movies, Culture
150111	*JayeCooley, Baddie___bey, ylreajd, LeanandCuisine, OModutle, RocFoster4, JeiMonroe, kodonism, BakingSodaYola, Thundercat*	METGala, Endgame, themasters, US, MetGala, TBT, 90s, dogs, pets, iPhoneVsHuaweiP30	International, Entertainment, Technology, Lifestyle
142499	*TravelVida, humorandanimals, F1, TravelPage, archpng, radnature, MercedesAMGF1, BestMovieLine, MercedesBenz, LetsbeAdventure*	art, photography, F1, debk, ebook, travel, porsche, Amazon, porsches, NowPlaying	Traveling, Technology, Hobbies
219563	*MontanaBlack, unge, NVIDIAGeForceDE, NetflixDE, Taddl, GermanLetsPlay, MaxAdlersson, Zombey, Paluten, FF_XIV_DE*	NintendoSwitch, ESC2019, SURO, Europawahl2019, Fortnite, Eurovision, GNTM, Rezo, Splatoon2, Artikel13	German, Entertainment, Politics, Gaming
224357	*KPWKM, rrrrrafla, asmolsushi, dauspozi, rlthingy, youngkessi, AfiqBushido, JKMHQ, latifborgiva, Nsyuhailarahmat*	SedangDiMainkan, WayV, 威神V, WeiShenV, UCL, 李永, TEN, SudirmanCup2019, WINWIN, DrakeCurse	Asian Pop-Culture
257645	*realDonaldTrump, bobcesca_go, marklevinshow, IvankaTrump, CNN, SirajAHashmi, nytimes, Twitter, _rshapiro, FLOTUS*	WWG1WGA, Weathercloud, Germany, Deutschland, MAGA, Election2020, Iran, Bayern, YesWeCan, Allemagne	International, Politics

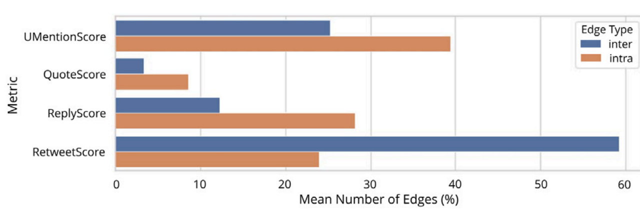

Fig. 7. Comparison of dominant interaction metrics within the communities.

These readers are, to a great extent, members of large communities. Traditional news providers, such as Spiegel, SZ, and Welt, spread in considerably more communities than newer providers. Therefore, they generated a greater reach with a broader audience. The massive audience also reflects itself in the lower PageRank of traditional media since many casual users are not well-connected in the network.

4.2.5. German Twitter Core Community

Finally, we explore the German Twitter Core Community (Core) (see Section 4.2.4) as it substantially determines the content we observe in the German-speaking Twitter community. Table 19 lists the most influential users within the Core by their PageRank percentile. At the top of the list, we find *Rezo*, the YouTube influencer, who was at the center of a political controversy surrounding the 2019 European Parliament election.[5] Besides one of the most discussed news topics revolving around *Rezo*, his profile also is an active, influential part of the Twitter community. By comparing his user degree with the degrees from news provider accounts, such as *Spiegel*, *Welt*, *FAZ*, etc., it is also apparent that more users interacted with him than with already established media profiles. Other political actors of the controversy, such as *Annegret Kramp-Karrenbauer* (@akk), and the official *CDU* profile (@CDU), are also present in the top user list. Users also utilized the Twitter profile of the *change.org* petition website (@ChangeGER) to attract attention to various topics throughout the Rezo controversy. We observe that most popular profiles are related to content contributing to political opinions. Furthermore, we found that news providers and politicians could establish widely popular Twitter profiles that stand at the top of the communities we discovered in our network. Therefore, we conclude that politicians adapted to the digital environment of Twitter and that German Twitter users show massive reactions towards them.

We analyzed the URL distribution in the Core and found that the user base is highly interested in external content from the News Group. 57% of all shared URLs contribute to political discussions. Additionally, 68% of the reaction-tweets relate to content from the News Group. Overall, we observed that 42% of the users in the Core discussed or shared news-related URLs. The observations suggest that the Core mainly discusses political content and consumes news media. This large-scale community indicates that active German Twitter users form a well-connected cluster rather than several smaller groups.

4.3. News discussion analysis

Besides information on news content, we are interested in the user behavior related to discussions (distinguished by the type of supported content). We augment our findings with information on community structures of the German-speaking Twitter community. Statistics on their activity, tweeting behaviors, and communication with other groups allow us to analyze group dynamics and -characteristics.

Tweeting behavior. Based on the tweeting behavior (see Table 17), we see that different communities exhibit diverse and partly contrasting tweeting practices. The willingness to communicate varies significantly between them. We identified two generic types we reference as *active*- and *passive groups*. For example, a high percentage of replies and quotes within the Core suggests that its users frequently engage in discussions (see Table 17 [left column]). Similar behavior is measurable within all communities of the active group. Groups related to politics show further emphasis on replies. On the other hand, passive communities mainly retweet ($\geq$ 60%). These figures indicate that their user bases mainly distribute content from other users.

Further statistics confirm these characteristic differences in tweeting behavior. Dissecting the interaction metric by its separate scores, we observe that, in most cases, one indicator dominates the others. For example, if user A frequently shares the contents of user B via retweets but only occasionally replies to them, the result will show a high retweet score and a low reply score. We consider the metric with the highest score as the dominant metric of an edge. These dominant metrics give us a more detailed view of the structure of communities. Table 17 (right column) gives an overview of the dominant metrics broken down by inter and intra-edges. We observe that active communities show a high percentage of user-to-user links dominated by replies (S_γ) and user mentions (S_M).

An essential trait of a sound community structure is not just isolated groups but users that have ties with users outside their community. Table 17 (middle column) details statistics on interconnections. We identify a perceptible difference between users of passive and active communities. Active communities tend to have a high share of interconnected users, suggesting a more active engagement in conversations than groups with a low amount of connected users.

Utilizing the Pearson correlation coefficient, we observe that the share of inter-connected users positively correlates with the ratio of replies (p = 0.63) and original tweets (p = 0.49) in communities. In contrast, retweet-heavy communities have fewer connections to other groups ($p = -0.6$).

[5] https://en.wikipedia.org/wiki/Rezo

Table 17
The 20 largest communities with their distribution of Tweets and additional information; μ and $\tilde{\mu}$ represent mean- and median values, respectively; UwE describes users with edges (to other communities).

ID	Users	Tweets	Per User		OT	RT	RP	QT	3rd	Inter-Edges		UwE	User Deg.		Dominance of Inter-Edges				Dominance of Intra-Edges			
															S_β	S_γ	S_τ	S_M	S_β	S_γ	S_τ	S_M
	#	#	(μ)	($\tilde{\mu}$)	%	%	%	%	%	#	(μ_U)	%	(μ)	($\tilde{\mu}$)	%	%	%	%	%	%	%	%
Core	816 677	47 737 955	67	4	24	37	36	8	10	9 436	1.15	34.75	45.59	3	34.07	22.06	3.52	40.35	39.07	21.32	2.18	37.43
109805	498 305	1 403 790	3	1	10	73	12	7	1	1 751	1.69	9.30	4.69	1	59.53	12.74	2.85	24.88	68.36	5.98	1.44	24.21
262453	381 148	1 236 799	3	1	13	67	18	3	3	1 560	1.86	14.76	4.37	1	64.51	16.02	2.58	16.88	87.28	6.36	0.96	5.40
34263	261 370	936 647	4	2	11	72	11	7	1	457	1.31	21.01	4.92	2	72.86	9.18	2.32	15.63	80.22	5.72	2.61	11.44
249774	261 057	440 439	2	1	7	86	4	7	1	1 786	1.88	17.56	2.87	1	91.59	2.39	0.93	5.10	92.28	3.07	1.79	2.87
150111	181 828	350 654	2	1	14	67	8	17	2	587	1.36	18.60	2.88	1	69.47	8.65	4.96	16.92	75.80	5.31	5.97	12.93
142499	179 695	1 376 486	9	1	35	53	10	3	21	1 049	2.86	24.49	6.55	1	56.89	10.13	2.63	30.34	62.58	3.85	1.05	32.52
219563	172 529	5 924 941	38	3	26	18	55	3	3	568	2.44	41.63	17.53	2	36.88	29.08	2.36	31.67	24.12	35.19	1.45	39.24
224357	165 974	342 187	2	1	10	73	10	13	3	596	1.53	15.52	3.36	1	77.25	9.38	2.28	11.09	72.73	10.72	2.95	13.60
257645	160 591	449 073	3	1	24	41	25	15	11	607	2.61	20.17	6.17	2	25.26	14.60	7.88	52.25	22.35	13.41	2.35	61.89
242038	149 786	276 990	2	1	11	76	6	13	1	381	1.42	14.16	2.91	1	85.18	3.54	1.40	9.88	79.55	6.13	3.53	10.79
143859	147 317	345 582	3	1	18	61	14	11	5	429	2.06	23.34	4.61	2	49.48	9.22	3.78	37.51	48.22	9.62	1.92	40.24
225111	135 624	319 732	3	1	13	64	13	18	1	506	1.58	20.73	3.42	1	72.53	7.58	3.28	16.60	63.11	11.74	4.98	20.17
96059	132 954	1 313 323	11	1	29	48	20	3	21	370	1.97	13.48	7.95	2	40.61	26.52	2.54	30.33	40.00	13.57	1.76	44.67
182077	132 654	1 663 479	13	2	25	32	40	6	7	784	3.83	32.75	11.69	2	31.77	22.47	2.19	43.57	29.62	22.00	1.74	46.64
34532	99 098	1 256 112	15	2	33	30	35	3	20	305	2.95	35.73	11.00	2	28.73	28.20	1.61	41.45	28.07	20.26	1.07	50.60
129034	97 746	343 890	4	2	13	72	8	8	2	246	1.94	24.79	5.60	2	69.24	8.66	2.60	19.49	75.62	4.29	2.96	17.13
195201	96 761	197 615	2	1	22	43	26	12	13	239	2.87	8.71	4.41	1	43.52	13.43	4.64	38.41	29.03	19.10	2.18	49.69
231865	81 920	298 713	4	1	14	66	16	6	3	210	2.21	7.28	4.90	1	49.07	22.29	3.69	24.95	64.92	7.86	2.95	24.28
32152	78 812	288 402	4	1	25	31	37	10	9	274	3.59	16.97	4.44	1	36.04	19.02	3.50	41.44	36.97	13.03	3.93	46.06

Table 18
Media influence in the community structure based on global PageRank percentiles and interconnectedness.

Content provider	Shared URL-Tweets			Reaction-Tweets		
	Com.	PRank	Deg.	Com.	PRank	Deg.
	#	∅	∅	#	∅	∅
spiegel.de	413	0.74	261	318	0.78	289
welt.de	320	0.75	330	300	0.77	300
bild.de	251	0.74	367	244	0.78	341
sueddeutsche.de	350	0.76	299	251	0.81	334
zeit.de	301	0.77	314	226	0.82	359
faz.net	247	0.78	352	194	0.82	359
focus.de	213	0.78	436	149	0.84	468
tagesschau.de	281	0.77	359	271	0.79	338
tagesspiegel.de	223	0.80	386	176	0.84	411
tichyseinblick.de	69	0.79	511	66	0.85	583
presseportal.de	134	0.82	546	82	0.86	590
journalistenwatch.com	53	0.82	585	58	0.85	615
taz.de	181	0.80	398	133	0.86	484
heise.de	226	0.75	314	159	0.81	381
n-tv.de	178	0.81	483	144	0.84	469
nzz.ch	203	0.77	368	140	0.83	453
handelsblatt.com	182	0.82	433	124	0.86	492
epochtimes.de	79	0.81	590	60	0.87	649
philosophia-perennis.com	63	0.81	586	62	0.85	631
zdf.de	223	0.80	415	177	0.84	421
derstandard.at	163	0.79	433	128	0.83	482
change.org	194	0.70	199	126	0.84	455
jungefreiheit.de	58	0.81	580	51	0.85	599
deutschlandfunk.de	139	0.83	492	118	0.88	544
wdr.de	148	0.84	482	114	0.87	534
bundestag.de	156	0.80	422	103	0.87	566
br.de	144	0.83	508	116	0.88	558
stern.de	160	0.81	515	137	0.83	519
ndr.de	151	0.84	518	124	0.88	560
rt.com	105	0.78	543	108	0.83	594

These findings suggest that communities with more inter-edges actively discuss the same contents as their adjacent communities. Further, active discussion culture seems to bring users from different communities together.

Finally, we observe that while passive communities related to Asian pop or entertainment and gaming show coherent activeness (mean-median ratio $2-3$), active groups related to German politics exhibit significantly different ratios (≥ 13). This discrepancy between a high mean value of tweets per user to a significantly smaller median indicates a small group of very active users within a community.

Communication patterns. We further observe peculiarities in internal- versus external communications. Fig. 7 depicts the average shares of dominant metric scores per community, separated by edge-type. The most striking difference indicates that retweets are more frequent between communities (Inter: 59%) than within communities (Intra: 24%). Furthermore, user mentions are the predominant type of interaction in communities (S_M (Intra): 39%), whereas retweets come in third after replies (S_γ (Intra): 28%). It suggests that discussions are the main factors for the forming of communities. In contrast, retweets dominate the connections between communities. We believe that users share content discovered outside of their community, supporting it via retweets. On

Table 19
Most influential users of the German Twitter Core Community.

Screen name	Name	Pagerank	Degree	Tweets #	3rd %	Label
@rezomusik	Rezo	0.999996	66 244	378	0	contro.
@CDU	CDU Deutschlands	0.999995	75 704	1 710	0	party
@akk	A. Kramp-Karrenbauer	0.999993	63 310	517	9	politician
@ChangeGER	Change.org DE	0.999991	20 939	312	0	activism
@SPIEGELONLINE	SPIEGEL ONLINE	0.999988	50 051	2 994	53	media
@welt	WELT	0.999986	45 410	11 424	99	media
@tagesschau	tagesschau	0.999982	52 230	2 816	87	media
@faznet	FAZ.NET	0.999978	36 622	4 810	78	media
@DB_Bahn	Deutsche Bahn	0.999977	17 508	12 462	100	info
@janboehm	Jan Böhmermann	0.999977	33 045	861	1	contro.
@BILD	BILD	0.999976	35 475	8 161	97	media
@DiePARTEI	Die PARTEI	0.999976	28 915	402	0	party
@KuehniKev	Kevin Kühnert	0.999975	39 372	289	0	politician
@Gronkh	GRONKH	0.999973	19 953	713	39	influencer
@zeitonline	ZEIT ONLINE	0.999972	36 751	3 585	95	media
@SZ	Süddeutsche Zeitung	0.999971	39 390	3 608	92	media
@sebastiankurz	Sebastian Kurz	0.999970	22 778	357	0	contro.
@nicosemsrott	Nico Semsrott	0.999970	33 749	270	0	politician
@spdde	SPD Parteivorstand	0.999969	36 288	4 988	57	party
@iBlali	Vik	0.999965	20 141	520	0	influencer

the other hand, it is rare for these users to comment on content from users outside their community via replies or quotes. Nonetheless, they reference users from other communities via user mentions (S_M (inter): 25%), showing a certain level of direct interaction.

4.4. Controversial users

Up to that point, we concentrated on content-related characteristics and behavior patterns. We complement our studies, exploring behavior patterns of users sharing controversial, anti-democratic content.

According to Section 3.3.3, we label users as either Controversial- or Non-Controversial Users. We detect 11 129 Controversial Users, identify their most influential members, and subsequently survey them to validate the group of Controversial Users. Thereby, we found several profiles from political personalities within the far-right ideological spectrum. These included well-known activists from the alt-right movement and German politicians from the right-wing party AfD. We also discovered authors from political blogs such as *philosophia-perennis.com*.

Related work [25–27] reported on political echo chambers from the extreme ends of the political spectrum. A common assumption regarding users within these chambers is that they only inform themselves based on a small and narrow set of information sources. McPherson et al. [56] reported this biased information consumption in social networks, called selective exposure. We focus on potential differences between controversial and non-controversial users and possible echo chambers revolving around anti-democratic content.

4.4.1. User base

Overall, we have a group of 11 129 users who support anti-democratic content. In the top 30 news providers on Twitter there are also 3 which spread anti-democratic content. *Epoch times*, supported by 6 900 users, *Journalistenwatch* supported by 6 471 users and *Philosophia perennis* supported by 6 408 users. The group of users supporting at least one of these 3 domains includes 10 694 accounts. 3 555 of which share articles from each of these 3 sources. Furthermore, it can be observed that a large part of these users also share articles from politically right-winged platforms that we do not consider to be extreme (*Tichy's Einblick*, *Junge Freiheit*), e.g. there are still 2 922 users who share articles from each of the 5 platforms (*Epoch Times*, *Philosophia perennis*, *Journalistenwatch*, *Tichy's Einblick* and *Junge Freiheit*.

In terms of responses to URL tweets from these providers, 12 809 users participated in the discussions (*Epoch Times* 7 542, *Philosophia perennis* 7 558, *Journalistenwatch* 7 616). Combined, this results in a group of 15 811 users who share or discuss these articles. Including *Tichy's Einblick* and *Junge Freiheit*, this figure grows to 22 334 with 19 043 users that responded to these URL tweets.

4.4.2. Tweeting behavior

Based on their PageRank (Mean PageRank: All 0.52/Non-Controversial 0.64/Controversial 0.76), Controversial Users are considerably well-connected in the network. The high reach of their tweets suggests that their overall influence is above average within the German Twitter user base. To understand how this influence manifests itself in the network, we study which hashtags they distribute.

Controversial Users produce a large share of political hashtags (see Def. 3.3.1). For example, the #AfD hashtag appears in 469 987 tweets shared by 49 883 users. While Controversial Users only make up for 15% (7 239) of these users, they generated 55% of these tweets. We made similar observations for most of the other tweets regarding political hashtags, such as #Merkel, #Islam, and #Flüchtlinge (eng.: refugees), and #Migranten (Engl.: immigrants). Despite their small numbers, Controversial Users, on average, distribute 42% of the tweets that contain political hashtags.

We further analyze the distribution and commenting behavior of Controversial Users. While these users prefer controversial information sources, they also share many articles from traditional news providers. In particular, articles from the large daily newspapers *Welt* and *Bild* (both conservative) attract a considerable attention from Controversial Users. Overall, 92% of the Controversial Users shared a traditional news provider at least once. They also use a wider variety of content providers (∅ 6) than Non-Controversial Users (∅ 3) to inform themselves.

So far, we only considered domains of external sources shared by Controversial Users. We extend our studies by exploring their reactions towards domains. We discovered that Controversial Users mainly react to traditional news sources. Articles from *Welt* caught the attention of many users in this group. Primarily, these users commented on political news articles that voice critical opinions about the *AfD* or reports about topics like immigration or political and climate activism. A closer look at related articles revealed several instances of comments on misinformation content (e.g., faked statistics) in favor of critical views about immigration. Moreover, Controversial Users fiercely commented against news articles that were generally positive on Islam and immigration. In contrast, extreme information sources received virtually no attention from Non-Controversial Users.

Our findings contradict the assumption that users with extreme views tend to form closed systems only reaffirming each other's beliefs. Based on the PageRank scores, the average Controversial User is more active than the average Non-Controversial User. Its average member achieves a higher reach in the German Twitter network than people with moderate political views. Most of their interactions with external

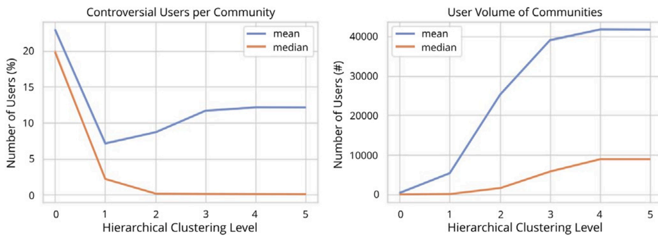

Fig. 8. Detailed statistics on the distribution of controversial users per iteration of the Louvain Method.

political content are responses to Tweets from Non-Controversial Users. They actively engage in many discussions and confront people with opposing views. Non-Controversial Users, on the other hand, tend to remain in their moderate area of political discussions, ignoring external content that supports extreme political ideologies.

4.4.3. Controversial communities

These findings contradict the notion of echo chambers. With high confidence, we can rule out epistemic bubbles. In general, however, they are no proof of the non-existence of echo chambers. We have to analyze controversial groups further. Do they form a type of echo chamber, which persistently discredits contrary political opinions (see Nguyen [28])?

In 2016, Zick et al. [57] reported a rising social acceptance for right-wing world views in Germany. This trend resulted in people expressing political opinions in public that would have been socially unacceptable before. So, while our results confirm that the average Controversial User does not withdraw into segregated groups reaffirming their political views, the question remains if this observation correlates with the development reported by Zick et al. [57]. We perform further studies to understand the diffusion of Controversial Users within the network. We leverage the hierarchy of our community structure. By examining each iteration of the Louvain method, we trace small user groups before they get merged into larger communities. Focusing on Controversial Users, we study the evolution of their memberships to communities through the hierarchy.

The first iteration places the Controversial Users into 1 734 communities (see Table 10). With further iterations, the number of communities drastically shrinks. It shows that the algorithm merges these user groups into larger, more general groups. After the last iteration, we observe 112 communities that include at least one Controversial User.

Fig. 8 gives an overview of the number of Controversial Users per community at different levels of Louvain clustering. We see the share of Controversial Users at the first level is relatively high. On average, 23% of the users in these communities are Controversial Users, which indicates denser controversial groups. However, we notice that most of these dense controversial groups are considerably small. User clusters at this resolution only reflect the communication between a few people and do not necessarily indicate political echo chambers but rather identify small-scale relationships between few users. Nevertheless, we also identified several user groups ranging from 20 to 80 members that solely shared extreme domains and hashtags. For example, several small communities mainly supported content from *anonymous.ru*, *pi-news.com*, or *philosophia-perennis.com*. They also showed little to no interest in traditional news media sources.

At successive levels, we register that most communities get merged with others. As a result, the median share of Controversial Users per community decreases. Interestingly, the mean share, after dropping on the second iteration, increases again on successive iteration. This effect continues with each hierarchical level. After the first iteration, most of the dense controversial groups we found at the bottom level are already part of the Core. The increasing mean share reflects the handful of controversial clusters that remain, for example, a network of 63 members that heavily support content from, e.g., *anonymousnews.ru*.

5. Discussion

The experiments provide extensive insights into the news consumption patterns of German Twitter users. In the following, we discuss the state of news consumption within the GTC.

General state. Interested in the share of news consumption within the GTC, we measured the exposure to news-related content. Leveraging shared external content, we observed that 25% of all users actively shared external sources. Regarding URL-tweets – posts that contained at least one URL to an external source – 41% were related to news content. Accounting only for unique URLs, we observe that news-related content makes up 29% of these. The discrepancy, a share of 29% unique URLs making up for 41% of all links, further emphasized the popularity of news-related content.

To refine our measurements on news consumption, we incorporated information about the community structures within the network. Studies on shared content within communities revealed that ~23% of the 25 572 identified communities, including 90.6% of all users, shared URL-tweets. ~28% of these communities shared news-related content. These 1 678 news-related communities (6.56% of all communities) produced 99% of all tweets within the network. Overall, 33% of all tweets in our data set supported or discussed news-related content. A similar picture emerged when analyzing the Core, the largest community within the GTC. With 57% of the URLs and 68% of Reaction-tweets related to news, 42% of its users shared or actively discussed such content.

Regarding the impact of news content, especially, the high number of replies w.r.t. tweets sharing and discussing such content suggests that users are willing to discuss or comment on others' content. The ratio of retweets of shared content (news related: 3.81; others: 2.15) reaffirmed this observation. Statistics on Reaction-tweets depicted a similar picture. News-related content and, especially, news providers triggered many Reaction-tweets. These figures suggest a high interest and participation in news consumption and political discussions. Only self-promotional profiles seemed to fall short of their intended goals, yielding minor to no effects concerning user engagement.

Regarding traditional news providers, the most popular outlets successfully established influential accounts within the GTC, with *Spiegel*, *Welt*, *Bild*, *Sueddeutsche*, *Zeit*, and *FAZ* at the top of all content providers within the GTC. Further, related German TV stations (e.g. *ZDF*, *WDR*, *BR*), related content (e.g. *Tagesschau*), and traditional news outlets from Switzerland (*NZZ*) and Austria (*derstandard.at*) were also part of the top 30 content providers. We also identified political blogs (e.g. *Tichyseinblick*, *Journalistenwatch*, *Epochtimes*, *Philosophia-perennis*) with a tendency towards tendentious to extreme content within the top 30. In this context, we also observed that the German political party *AfD* was highly active on social media. Regarding shared links from Facebook (6 out of 10) and YouTube (3 out of 10), *AfD*-related topics dominated this content.

Besides news providers directly operating on Twitter, only *BTS* and *Rezo* were able to generate reach with links from *YouTube*. Here, one event showed the impact the *BTS*-community can have on networks. A single link was shared 284 980 times by 282 903 users. In comparison, the most shared news provider, *Spiegel*, reached a total tweet-count of 344 946 with 47 805 different links. However, *BTS* and *Rezo* are the exception. Regarding other content, links from YouTube, Instagram, and Facebook failed to generate reach and impact. Here, users shared most of the content from Instagram and Facebook via third-party apps, indicating a more passive Twitter use.

Taking community structures into account, influential nodes, besides traditional news providers, were politicians, the political parties *CDU*, *Die PARTEI* and *SPD*, streamers/influencers, and accounts in conjunction with controversial topics. The only two other accounts in the 20 most influential accounts were the activism platform *change.org* and the Twitter account of *Deutsche Bahn*, the national railway company of Germany. The data revealed that young politicians (@KuehniKev,

@nicosemsrott) established widely popular Twitter profiles. Further, political controversies seemed to have an immediate and significant impact on the network structure. For example, the controversy surrounding *Rezo* significantly influenced the reach and visibility of not only himself but also of profiles affected by the event (see, e.g., *@akk*, *@CDU*). It demonstrates the impact a political actor can achieve on short notice.

Finally, we observed that activeness is a characteristic of news-related communities. Others, such as entertainment-, gaming-, or lifestyle-related ones, showed a high share of retweets. News-related communities, however, exhibited a more dynamic behavior via replies and quotes. We also observed that large communities include users from the whole political spectrum.

Controversial news-content. To classify observations on controversial news-content, we need to look at related work. Bor and Petersen [58] examined the question of why online discussions seem more hostile than their offline counterparts. They examined eight studies using cross-national surveys and behavioral studies and concluded that it is not that people are more hostile online, but that hostile people gain greater visibility online. Additionally, other studies report that emotion triggering posts [59], especially posts about political opponents are substantially more likely to be shared [60]. Combined, these effect seems to be amplified by the fact that moderate users turn away from discussions because of this hostile behavior [44]. This inevitably leads to the behavior of the few receiving a disproportionate amount of attention. In the U.S., this seems to be compounded by the fact that the most extreme left- and right winged political groups not only attack users with opposing views, but are particularly hostile to moderates who espouse their beliefs [44,61,62]. Hawkins et al. [44]"Those who express sympathy for the views of opposing groups may experience backlash from their own cohort". This behavior undermines discussion between people with different opinions and even causes social media to have a detrimental effect on democratic societies [63].

Our work now sheds light on the situation in German-speaking countries. Established news providers dominate news-related content within the GTC. Nonetheless, actors spreading and supporting controversial opinions are also part of the landscape. We observed striking differences in the supporting patterns of different news types. While moderate news was widely shared and discussed, users supported tendentious news sources mainly via retweets. Supporters of extreme political content use the Reply-function (22–24%) to inject their content into discussions. Moderate users, however, mostly ignore it.

We extended our research on controversial news content, focusing on providers and users supporting tendentious to extreme sources. Here, two strongly varying pictures emerged. On the one hand, content providers that distribute tendentious to extreme political content play only minor roles in the network (see Table 13). On the other hand, their supporters are highly active and noticeably well connected.

At first glance, this high frequency of interactions with various users contradicts the assumption that people with more extreme political ideologies tend to form echo chambers [7]. However, similar to Zick et al. [57], our findings suggest the existence of a more self-confident form of echo chambers. By dissecting the different layers of the network partition, small coherent groups with selective exposure to extreme political content emerged. Interestingly, these groups became part of larger clusters that predominantly engaged in discussions of moderate political content. Taking their high activity, hashtag usage, content, and shared URLs into account, a picture similar to the findings in Hawkins et al. [44],Bor and Petersen [58] emerged. A minor group of extreme users – formed according to standard echo chambers – spread out to aggressively support their opinions in public. From a group – repellent to opposing views and reassuring in their political positions – these users evolved to highly active members of larger communities.

These users drastically increased their reach and visibility. While popular domains in the Top 30 that share anti-democratic content only have roughly ≈6 500 supporters (combined: 10 694) and an active audience of ≈7 500 users, *Journalistenwatch* (Position in Top30: 12th with 95 225 Tweets supporting and 84 145 Tweets discussing the content), *Epochtimes* (18th: 73 086 / 68 923), and *Philosophia perennis* (19th: 72 926 / 49 726) are among the 30 most shared news providers in the GTC.

For example, 9 088 articles of the renowned newspaper *Süddeutsche Zeitung* were discussed by 53 591 users in 190 727 tweets (*Tweets/audience*: 3.56, *Tweets/article*: 20.99), while only 721 articles of the anti-democratic domain *Philosophia perennis* were discussed by 7 558 users in 49 726 tweets (*T/audience*: 6.58, *T/article*: 68.97). So, 7-times more people discussed articles of the moderate outlet in comparison to articles of the anti-democratic domain. The moderate discussions, however, only generated 3.8x more tweets with only 2x the number of retweets involved in the anti-democratic discussions. Here, 68% of the 'discussions' were in form of retweets.

In summary, we conclude that a similar behavior from users of the extreme ends of the political spectrum as reported in Hawkins et al. [44] can be observed in the GTC. The average controversial user has a high PageRank, i.e., a user's profile connects to other well-connected users within the GTC. Interestingly, however, it seems that these users are largely ignored in discussions by the moderate majority of users in the GTC.

Due to missing data of previous years, we could not study potential developments, e.g., if it correlates to the rise of social acceptance of their opinions [57].

6. Limitations

We reported exhaustive studies on the influence and impact of anti-democratic news content. To cope with a large data set, we formulated several assumptions. Thereby, we accepted certain limitations of our approach.

Content understanding. We based our study on a large data sample. Thereby, we decided to rely on automated methods for content understanding. Studying the content discussed on Twitter via shared external content seems a rough estimate in the first place. However, curated third-party services significantly reduce the complexity of content understanding. Looking at a handful of domains to understand an FG and, thereby, thousands of articles/domains helped us cope with the sheer amount of data. Also, due to the restrictions on tweet length, URLs offer themselves an easy way to share opinions.

Statistics on our data set support and confirm our abstraction approach. Alone 1/3 of all tweets discussed news-content. Including other discussed content, the method allows understanding large parts of discussed topics.

Data collection. We collected users active during the collection phase. Therefore, we missed all inactive users, even if these users passively consumed content on Twitter. Follower-information would have provided data on passive users (having other drawbacks). The information would also have allowed for more detailed approximations of reach and impact of content. However, concentrating on a virtually complete snapshot of the targeted community made it impossible to collect this information (request limitations).

Promotional profiles. Finally, our crudest approximation regards promotional profiles. To rely on voluntarily provided information from the relevant account carries some risks. Especially the striking difference between traditional news providers (where we identified plenty of promotional profiles) and tendentious to extreme news-content providers (almost none) needs further investigations. To mitigate these uncertainties, one could use shared external content information.

Further research on the detection of automated accounts is also needed. We decided to ignore the noise introduced by bots, because recent reports question current detection solutions. According to Majó-Vázquez et al. [38], e.g., accounts we focused in our research are especially prone to get suspended due to their behavior rather than bot activities.

Controversial users. Controversial users are almost surely correctly labeled. To ensure this, we only labeled users as *controversial* that actively shared an article from extreme, anti-democratic domains. This probably leads to the fact that we have an uncertainty in the group of non-controversial users. However, we argue that the imprecision introduced in this way has a smaller impact (arguably none) because it affects by far the larger group of users to a much smaller extent.

7. Conclusion

This work focused on Twitter's German-speaking user base and their behavior. We emphasized external information sources contributing to the forming of political opinions. The goal was to estimate the influence of anti-democratic political information on the German community. We captured the Twitter traffic of German-speaking users over two months during the 2019 European Parliament election. By utilizing the Twitter Streaming API for our systematic collection approach, we obtained a representative snapshot of the German-speaking Twitter community, comprised of 77M tweets from 6.9M users.

Evaluations yielded detailed insight into the daily and weekly routines of the German Twitter population. In particular, we found that political events, such as several controversies and the election, lead to significant activity increases. To further study the news consumption of users, we categorized external content. The automated categorization successfully assigned categories to 98% of the URL-sharing tweets. We believe that such an approach provides a powerful tool for identifying meta-information in large-scale networks.

Information contributing to political opinion-forming received the most reactions from the German user base. The most prominent figures include traditional news media, official governmental sites, and political blogs within the far-right political spectrum. Due to the election period, official governmental information providers also received much attention and were referred to by users on the platform. Traditional content providers also put much effort into creating sophisticated promotional networks within Twitter. News-feed accounts and journalists contributed to the distribution of articles and became involved in political discourse.

Successive partitions of the clustering algorithm provided further evidence on the influence of political content on community structures. Users in dense clusters mainly interact via commenting on each other's content or referencing each other via user mentions. This active discussion culture on political content is the driving force behind the formation of several large-scale communities in our network. In the center of the interactions stands the German Twitter Core Community, which produced 62% of the tweets in our corpus and included 814k users. The group is highly active and evolves around political events and personalities. The most-read traditional media providers established profiles within the Core. Based on their PageRank, these users are highly influential in the GTC. Politicians, political parties, and controversial personalities are the most influential nodes in the communities. The finding highlights how well-connected political actors are. They reach a broad audience.

With the acquired background information, we focused on the research questions. We studied the scale and influence of anti-democratic news content. Thus, we defined the groups of controversial- and non-controversial users. Comparing their tweet behavior, we found striking differences. Mostly, news providers received significant attention from the user base and contributed to a lively discussion culture. In contrast, people who consumed tendentious to extreme politically opinionated blogs were overly supportive, but users discussed their content far less. These small blogs, supporting extreme political ideologies, had a small but loyal user base that distributed their content extensively via retweets in the network.

We observed that most communities include users from all over the political spectrum. Thus, we could not confirm the existence of massive networks of ideologically segregated user groups (cf. Boutyline and Willer [7]). However, similar to the study by Zick et al. [57], the data revealed that members of previously existing echo chambers started to support their opinions in discussions with dissenters. While our results provided evidence that small-sized user clusters, supporting extreme views, exist in the Twitter network, most of these communities became part of large-scale clusters. Similar to findings reported in the 'Hidden Tribes' study [44], these politically motivated controversial users are overly active on Twitter. Despite their small size (overall 11 129 users), they generate large amounts of tweets. Overall, people with extreme political views are well-connected and frequently engage in discussions with users that share moderate information sources. However, information on used hashtags suggests that these users propagate their opinions rather than discuss topics. Due to their high activity, this small group of users is overly influential and visible in the GTC.

These findings add to a growing body of literature on political polarization and the forming of echo chambers. We devised an innovative strategy to evaluate how small-sized political clusters become part of large-scale communities and used Twitter data to provide meaning to these structures. Based on our data, we believe that the German Twitter user base consumes daily news and is highly interested in interacting with political personalities directly on the platform. The most active communities on German Twitter evolve around these political figures and events. Meta-data and behavioral patterns of controversial users revealed the development of a new self-assured form of echo chambers.

Our work is a first step towards enhancing our understanding of the German Twitter population. We highly suggest that researchers apply similar methods to conduct their studies on large-scale snapshots rather than small network samples.

CRediT authorship contribution statement

Jan Ludwig Reubold: Conceptualization, Methodology, Resources, Writing, Formal analysis, Investigation. **Stephan Escher:** Conceptualization, Methodology, Resources, Supervision, Writing, Investigation. **Johannes Pflugmacher:** Methodology, Software, Validation, Formal analysis, Investigation. **Thorsten Strufe:** Conceptualization, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work in part was supported by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany's Excellence Strategy – EXC 2050/1 – Project ID 390696704 – Cluster of Excellence "Centre for Tactile Internet with Human-in-the-Loop" (CeTI) of Technische Universität Dresden and by funding of the Helmholtz Association, Germany through the KASTEL Security Research Labs (HGF Topic 46.23).

Appendix A. Functional groups

Table 5 details the distribution volumes of the Top 10 FGs and their categories. While the users generate most of their traffic in the *Information/Communica-tion* FG (47% tweets), the FG with the maximum user base is *Entertainment/Culture*. The majority of the users of this FG are interested in multimedia content, such as videos and photos (Streaming Media: 36%; Media Sharing: 33%).

Most of the content of the *Information/Communication* FG is related to news (General News: 32%) and personal blogs (Blogs/Wiki: 10%), mainly consisting of content from online news media and personalized political websites. Based on the high number of retweets in this

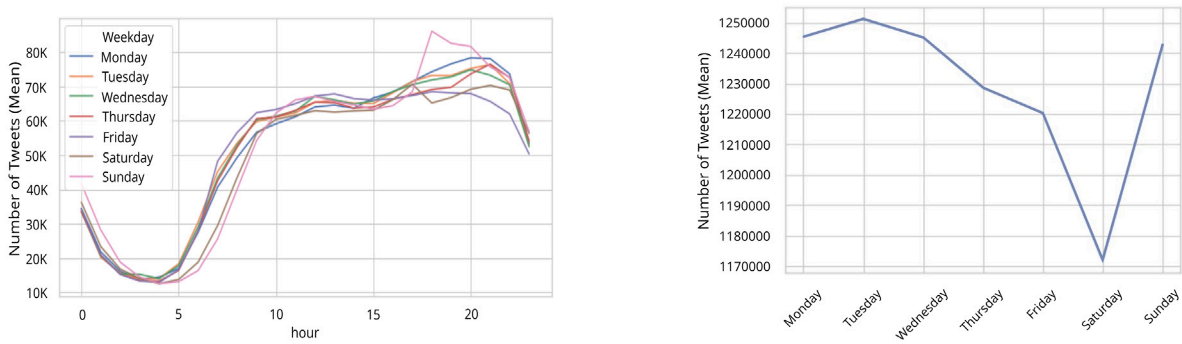

(a) Twitter activities over the course of every weekday.

(b) Twitter activities over the course of a week.

Fig. 9. Tweets over time.

group (52%), news and blog content seems to be well-received by the user base. We observed the same popularity of political domains in the FG *Society/Education/Religion*, comprised of even more elaborate political content. Most of the URLs captured during our data collection are from domains within the *Information/Communication* FG, which means a high number of different news articles are generated and distributed on Twitter. Despite this variety of articles, the Twitter community still reacts to these links by spreading them via retweets and replies. In contrast to news and political content, lifestyle-related content (FG: *Lifestyle*) results in fewer retweets (34%), which indicates a lower acceptance by the German community. An exception is the category Controversial Opinions, which includes domains that share highly opinionated political content (e.g., journalistenwatch.com, philosophia-perennis.de, pi-news.net). McAfee's TrustedSource grouped *Controversial Opinions* into the FG *Lifestyle*. The number of retweets in this category is 70%, further supporting the assumption that political content on Twitter is widely distributed and acknowledged.

The FGs with the most original tweets are related to marketing campaigns (*Purchasing*: 75%), business advertising (*Business/Service*: 71%), and online technologies (*Information Technologies*: 69%). Third-party services generate a majority of these tweets. Therefore, we assume that most of these domains conduct an automated distribution of their products. The lack of retweets within the respective categories indicates that this distribution approach is not overly effective in the German Twitter community.

Further, the share of spam and inappropriate content is relatively small (overall 4% to 7%). Just a tiny margin of users is involved in the distribution process. Spam URLs found in our data were mainly shared via original tweets (97%) and distributed via third-party services (92%), which suggests automated distribution in the context of spam and marketing. Based on the low number of retweets, users recognize spam content and do not distribute these any further in the network.

Appendix B. Tweets over time

The volume of daily captured tweets varies from 1M to 1.6M messages with an average of 1.2M. By examining the average collection of tweets by weekdays, we observed that German-speaking Twitter users were more active from Sunday to Tuesday and had a decreasing interest in Twitter from Wednesday to Saturday, with the lowest activity on Saturdays (see Fig. 9(b)). The overall daily usage (see Fig. 9(a)) is moderate in the morning, increases during after-work hours, and drops to its lowest point at night between 1 am and 5 am. At the weekend, Twitter usage naturally starts a few hours later in the morning. The oddly shaped peak on Sunday evenings results from high volumes of tweets during the night of the 2019 European Parliament election. The daily Twitter activities match Central European Time and the working schedule of people from Germany and Austria.

Appendix C. Captured events

We use hashtags to give an overview of relevant discussions in the network. The two-month time frame is intended to ensure that the data set is as diverse as possible. Below, we explore the relevant topics based on the most popular hashtags in our data set. Table 4 gives an overview of the busiest days on Twitter and the occurring Hashtags.

Twitter users were most active at the end of the campaign for the 2019 European parliament elections (May 26). All the top hashtags shared on Twitter during this period can be attributed to the election and the associated debate about the election results. The *Christian Democratic Union* (CDU) leads in both the hashtag ranking and the actual election (28.9%). The controversial *Alternative for Germany* (AfD) party follows next (#AfD), although it only came fourth in the election with a total of 11%. In comparison, the *Bündnis 90/Die Grünen* party, which came in second in the election with almost twice as many votes (20.5%), only appears in 11th place in the hashtag ranking. The *Sozialdemokratische Partei Deutschlands* (SPD) landed in third place in the election and also managed to attract more attention on Twitter (#SPD) than *Die Grünen*.

However, the popularity of the hashtag #CDU could also be a side effect of Annegret Kramp-Karrenbauer's (#AKK) controversial comments on the political comments of Youtube influencer Rezo (#Rezo), which triggered a general discussion about censorship in online forums (#Censorship). Looking at the context of the hashtags, one concludes that the popularity of the hashtags is the result of lively discussions rather than a reflection of political party affiliation. The beginning of the Rezo controversy can be seen in the spike in tweet volume on May 22 and 23. Youtube influencer Rezo (#Rezo) posted a viral video (#RezoVideo) to express his concerns about the CDU's political course. The video received widespread media coverage and led to a reaction video (never published) by CDU politician Philipp Amthor (#Amthor).

Another political controversy occurred on May 18. Austrian politician Heinz-Christian Strache (#Strache) was the main character of a compromising video (#StracheVideo) that caused the Austrian government coalition to collapse.

Based on the popular hashtags, we see that a high number of political topics are discussed. In addition, previously announced political campaigns on Twitter were also able to generate high volumes of tweets. The hashtag #NichtOhneMeinKopftuch was the most dominant hashtag on June 2, with 124 218 tweets. For comparison, the second most shared hashtag that day was mentioned in only 17 991 tweets.

In addition to political events, pop culture events also dominate Twitter (e.g. #GNTM, #ESC2019). There are also some non-German hashtags that refer to a Korean pop band called BTS, which reached high rankings in the music charts in Germany for several weeks. During our data collection, they released several singles and generated trending hashtags. Most popular events also dominated the news in Germany during the data collection period. Therefore, we can conclude that our data collection correctly collects German-language tweets.

Appendix D. External media sources

We discovered 1.4M tweets that shared 374k distinct *YouTube-URLs*. While the number of shared *Instagram-URLs* (520k) is only a third of the distributed *YouTube-URLs*, they contain a similar number of distinct *URLs* (370k). We observe the same when looking at content from *Facebook*. Regarding the type of media shared via these platforms *YouTube links* seldom contained other content than video links (97%). These videos originated from 97k *YouTube channels*. Via *Instagram* the most common shared media types are images (71%) followed by *PostPages* (12%), which also contain multimedia content and *profile pages* (10%). The content from *Facebook-links* is mainly textual (post: 53%; story: 13%) and less multimedia-based (photo: 10%; video: 8%). There are only a few events and groups shared within our corpus.

Links disguised by users via link shortening services make up 14% of URLs shared on Twitter. Resolving these links, we discovered that news providers and bloggers use marketing services to distribute their content in an automated manner. Another large share of disguised URLs is related to inappropriate content, such as pornography.

YouTube. Overall, the data set contains 1 402 441 tweets that shared 374 414 distinct YouTube-URLs originating from 97 680 different YouTube Channels. The content of shared videos varies from music, gaming, and political opinions to educational content (see Table 8). We identified single videos accounting for large chunks of the YouTube links on Twitter. For example, a newly released single of a Korean pop band (BTS) or a video of a channel called *Rezo* belonging to a person who was at the center of a political controversy surrounding the 2019 European Parliament election. He published a video with the title "Die Zerstörung der CDU" (Engl.: the destruction of the CDU) that went viral, expressing concern regarding the political course of the *CDU*. In general, there is only a small number of frequently shared content providers from YouTube (see Table 9). Half of these Channels are related to political topics. Moreover, they show a specific political affiliation. Channels belonging to the right-wing political party AfD are shared more often than channels of any other party. This observation indicates a high activity during their election campaign and shows a trend towards utilizing multimedia content to reach a broader spectrum of users.

Instagram. Although the number of shared Instagram-URLs (520 466) amounts to only a third of the distributed YouTube-URLs, they show a similar number of distinct URLs (370 510). Most of the content shared via Instagram links are images (71%), direct posts (12%), which also contain multimedia content, and profile pages (10%). Overall, shared Instagram links are mostly apolitical and dominated by profiles from the entertainment industry.

Facebook. The distribution ratio of Facebook links shows a high similarity to Instagram shares. In total 454 128 Facebook-URLs were distributed, with 292 316 distinct destinations, originating from 96 221 different Facebook profiles. By looking at the most shared Facebook profiles (see Table 9), we observe a relatively small user base that only supports a handful of Facebook pages or profiles, with a low distribution factor. We notice, however, that most Facebook profiles are politically motivated and shifted towards the right-wing party AfD. One exception to this rule is a frequently shared page that directly opposes said party (@GegenDieAfD).

Overall, the top content providers from YouTube and Facebook are mostly related to political parties and activism.

References

[1] S. Engesser, E. Humprecht, Frequency or skillfulness: How professional news media use Twitter in five western countries, J. Stud. (2015).

[2] S. Stier, A. Bleier, H. Lietz, M. Strohmaier, Election campaigning on social media: Politicians, audiences, and the mediation of political communication on facebook and Twitter, Political Commun. (2018).

[3] S. Hölig, U. Hasebrink, Ergebnisse für deutschland, Reuters Institute Digital News Report 2019, 2018.

[4] S. Zannettou, M. Sirivianos, J. Blackburn, N. Kourtellis, The web of false information: Rumors, fake news, hoaxes, clickbait, and various other shenanigans, J. Data Inf. Qual. (2019).

[5] N. Ernst, S. Engesser, F. Büchel, S. Blassnig, F. Esser, Extreme parties and populism: An analysis of facebook and Twitter across six countries, Inf. Commun. Soc. (2017).

[6] P. Barberá, J.T. Jost, J. Nagler, J.A. Tucker, R. Bonneau, Tweeting from left to right: Is online political communication more than an echo chamber? Psychol. Sci. (2015).

[7] A. Boutyline, R. Willer, The social structure of political echo chambers: Variation in ideological homophily in online networks, Political Psychol. (2017).

[8] E. Colleoni, A. Rozza, A. Arvidsson, Echo chamber or public sphere? Predicting political orientation and measuring political homophily in Twitter using big data, J. Commun. (2014).

[9] P. Boczkowski, E. Mitchelstein, M. Matassi, News comes across when i'm in a moment of leisure: Understanding the practices of incidental news consumption on social media, New Media Soc. (2018).

[10] A. Ardèvol-Abreu, T. Diehl, H. Gil de Zúñiga, Antecedents of internal political efficacy incidental news exposure online and the mediating role of political discussion, Politics (2019).

[11] S. Wineburg, S. McGrew, J. Breakstone, T. Ortega, Evaluating information: The cornerstone of civic online reasoning, Stanf. Digital Repository (2016).

[12] D. Lazer, A. Pentland, L. Adamic, S. Aral, A.-L. Barabási, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, et al., Computational social science, Science (2009).

[13] S. Roos, B. Schiller, S. Hacker, T.S. trufe, Measuring freenet in the wild: censorship-resilience under observation, in: PETS, 2014.

[14] H. Kwak, C. Lee, H. Park, S. Moon, What is Twitter, a social network or a news media? in: WebConf, 2010.

[15] A. Pal, S. Counts, Identifying topical authorities in microblogs, in: Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, 2011.

[16] R. Cohen, D. Ruths, Classifying political orientation on Twitter: It's not easy!, in: Proceedings of the International AAAI Conference on Web and Social Media, 2013.

[17] Y. Genc, Y. Sakamoto, J.V. Nickerson, Discovering context: Classifying tweets through a semantic transform based on wikipedia, in: FAC, 2011.

[18] C. Budak, A. Kannan, R. Agrawal, J. Pedersen, Inferring user interests from microblogs, in: Proceedings of the International AAAI Conference on Web and Social Media, 2014.

[19] G. Di Tommaso, S. Faralli, G. Stilo, P. Velardi, Wiki-MID: A very large multi-domain interests dataset of Twitter users with mappings to wikipedia, in: International Semantic Web Conference, 2018.

[20] S. Faralli, G. Stilo, P. Velardi, Large scale homophily analysis in Twitter using a twixonomy, in: IJCAI, 2015.

[21] I. Himelboim, M. Smith, B. Shneiderman, Tweeting apart: Applying network analysis to detect selective exposure clusters in Twitter, Commun. Methods Measures (2013).

[22] L. Jin, Y. Chen, T. Wang, P. Hui, A.V. Vasilakos, Understanding user behavior in online social networks: A survey, IEEE Commun. Mag. (2013).

[23] A. Schulz, B. Schmidt, T. Strufe, Small-scale incident detection based on microposts, in: ACM Hypertext and Social Media (Ht), 2015.

[24] C. Wilson, B. Boe, A. Sala, K.P. Puttaswamy, B.Y. Zhao, User interactions in social networks and their implications, in: Proceedings of the 4th ACM European Conference on Computer Systems, EuroSys '09, 2009.

[25] M.D. Conover, J. Ratkiewicz, M. Francisco, B. Gonçalves, F. Menczer, A. Flammini, Political polarization on Twitter, in: Proceedings of the International AAAI Conference on Web and Social Media, 2011.

[26] A. Boutet, H. Kim, E. Yoneki, What's in Twitter, I know what parties are popular and who you are supporting now!, Soc. Netw. Anal. Min. (2013).

[27] A. Feller, M. Kuhnert, T.O. Sprenger, I.M. Welpe, Divided they tweet: The network structure of political microbloggers and discussion topics, in: Proceedings of the International AAAI Conference on Web and Social Media, 2011.

[28] C.T. Nguyen, Echo chambers and epistemic bubbles, Episteme (2020).

[29] C. Orellana-Rodriguez, D. Greene, M. Keane, Spreading one's tweets: How can journalists gain attention for their tweeted news? J. Web Sci. (2017).

[30] Reuters, Reporting from the internet and using social media, 2013, Reuters, Thomson.

[31] A. France-Presse, AFP updates guidelines on using social media, AFP Newsletter [Online] (2013).

[32] Associated Press, Social media guidelines for AP employees, 2013.

[33] T. Scheffler, A German Twitter snapshot, in: LREC, 2014.

[34] D. Goldhahn, T. Eckart, U. Quasthoff, Building large monolingual dictionaries at the leipzig corpora collection: From 100 to 200 languages, in: LREC, 2012.

[35] F. Morstatter, Y. Shao, A. Galstyan, S. Karunasekera, From alt-right to alt-rechts: Twitter analysis of the 2017 german federal election, in: Companion Proceedings of the the Web Conference 2018, 2018.

[36] B. Pavliy, J. Lewis, The performance of Twitter's language detection algorithm and google's compact language detector on language detection in Ukrainian and Russian tweets, Bull. Toyama Univ. Int. Stud. (2016).
[37] J. Pflugmacher, S. Escher, J. Reubold, T. Strufe, The German-speaking Twitter community reference data set, in: INFOCOM/HotPOST, 2020.
[38] S. Majó-Vázquez, M. Congosto, T. Nicholls, R.K. Nielsen, The role of suspended accounts in political discussion on social media: Analysis of the 2017 French, UK and German elections, Soc. Media+ Society (2021).
[39] J. Echeverria, E. De Cristofaro, N. Kourtellis, I. Leontiadis, G. Stringhini, S. Zhou, LOBO: Evaluation of generalization deficiencies in Twitter bot classifiers, in: Proceedings of the 34th Annual Computer Security Applications Conference, 2018.
[40] A. Rauchfleisch, J. Kaiser, The False Positive Problem of Automatic Bot Detection in Social Science Research, Berkman Klein Center Research Publication, 2020.
[41] F. Riquelme, P. González-Cantergiani, Measuring user influence on Twitter: A survey, Inf. Process. Manage. (2016).
[42] M. Cha, H. Haddadi, F. Benevenuto, K.P. Gummadi, Measuring user influence in Twitter: The million follower fallacy, in: Proceedings of the International AAAI Conference on Web and Social Media, 2010.
[43] T.F. Van de Mortel, et al., Faking it: Social desirability response bias in self-report research, Australian J. Adv. Nurs. (2008).
[44] S. Hawkins, D. Yudkin, M. Juan-Torres, T. Dixon, Hidden tribes: A study of America's polarized landscape, 2019.
[45] P. Metaxas, E. Mustafaraj, K. Wong, L. Zeng, M. O'Keefe, S. Finn, What do retweets indicate? Results from user survey and meta-review of research, in: Proceedings of the International AAAI Conference on Web and Social Media, 2015.
[46] G. Gadek, A. Pauchet, N. Malandain, K. Khelif, L. Vercouter, S. Brunessaux, Topical cohesion of communities on Twitter, Procedia Comput. Sci. (2017).
[47] V.D. Blondel, J.-L. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks, J. Stat. Mech. Theory Exp. (2008).
[48] S. Fortunato, M. Barthelemy, Resolution limit in community detection, Proc. Natl. Acad. Sci. USA (2007).
[49] J. Reichardt, S. Bornholdt, Partitioning and modularity of graphs with arbitrary degree distribution, Phys. Rev. E (2007).
[50] Y. Wang, Forcing a breakdown: Establishing the limits of community detection algorithms, in: Sunbelt, 2012.
[51] M.E. Newman, Fast algorithm for detecting community structure in networks, Phys. Rev. E (2004).
[52] A. Clauset, M.E. Newman, C. Moore, Finding community structure in very large networks, Phys. Rev. E (2004).
[53] A. Lancichinetti, M. Kivelä, J. Saramäki, S. Fortunato, Characterizing the community structure of complex networks, PLoS One (2010).
[54] A. Lancichinetti, S. Fortunato, Erratum: Community detection algorithms: A comparative analysis, Phys. Rev. E (2014).
[55] S. Weischenberg, M. Malik, A. Scholl, Journalismus in Deutschland 2005, Media Perspektiven (2006).
[56] M. McPherson, L. Smith-Lovin, J.M. Cook, Birds of a feather: Homophily in social networks, Annu. Rev. Sociol. (2001).
[57] A. Zick, B. Küpper, D. Krause, Gespaltene Mitte–Feindselige Zustände, Rechtsextreme Einstellungen in Deutschland (2016).
[58] A. Bor, M.B. Petersen, The psychology of online political hostility: A comprehensive, cross-national test of the mismatch hypothesis, Am. Political Sci. Rev. (2022).
[59] W.J. Brady, J.A. Wills, J.T. Jost, J.A. Tucker, J.J.V. Bavel, Emotion shapes the diffusion of moralized content in social networks, Proc. Natl. Acad. Sci. (2017).
[60] S. Rathje, J.J.V. Bavel, S. van der Linden, Out-group animosity drives engagement on social media, Proc. Natl. Acad. Sci. (2021).
[61] K. Garimella, G. De Francisci Morales, A. Gionis, M. Mathioudakis, Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship, in: Proceedings of the 2018 World Wide Web Conference, 2018.
[62] J.T. Moss, P.J. O'Connor, Political correctness and the alt-right: The development of extreme political attitudes, PLoS One (2020).
[63] P. Lorenz-Spreen, L. Oswald, S. Lewandowsky, R. Hertwig, Digital media and democracy: a systematic review of causal and correlational evidence worldwide, 2021.