Opinion-based Homogeneity on YouTube

Combining Sentiment and Social Network Analysis

Daniel Röchert, German Neubaum, Björn Ross, Florian Brachten, Stefan Stieglitz

University of Duisburg-Essen

Abstract

When addressing public concerns such as the existence of politically like-minded communication spaces in social media, analyses of complex political discourses are met with increasing methodological challenges to process communication data properly. To address the extent of political like-mindedness in online communication, we argue that it is necessary to focus not only on ideological homogeneity in online environments, but also on the extent to which specific political questions are discussed in a uniform manner. This study proposes an innovative combination of computational methods, including natural language processing and social network analysis, that serves as a model for future research examining the evolution of opinion climates in online networks. Data were gathered on YouTube, enabling the assessment of users' expressed opinions on three political issues (i.e., adoption rights for same-sex couples, headscarf rights, and climate change). Challenging widely held assumptions on discursive homogeneity online, the results provide evidence for a moderate level of connections between dissimilar YouTube comments but few connections between agreeing comments. The findings are discussed in light of current computational communication research and the vigorous debate on the prevalence of like-mindedness in online networks.

Keywords: machine learning, echo chamber, social network analysis, computational science, opinion-based homogeneity

Opinion-based Homogeneity on YouTube

Combining Sentiment and Social Network Analysis

Social media such as YouTube, Facebook, or Twitter have fundamentally changed people's political communication by offering the opportunity to exchange opinions across time and geographical barriers. At the same time, there are risks associated with the use of social media, such as being exposed to manipulative agents like social bots, the viral spread of misinformation, or the formation of echo chambers, i.e., online spaces in which users exclusively encounter information and opinions in line with their own.

According to current research, these risks could (a) undermine the heterogeneity of opinion climates (Graham, 2015), (b) narrow (political) world views and even convey distorted pictures of public opinion to individual users (Neubaum & Krämer, 2017), and (c) foster a polarization of viewpoints and fragmentation of society (Sunstein, 2017). Empirical studies using computational methods (i.e. network analyses) have found that users in networks such as Twitter indeed move in ideologically homogeneous clusters, but are still confronted time and again with information and opinions divergent from their own (e.g., Bakshy, Messing, & Adamic, 2015; Guo, Rohde, & Wu, 2018) which, in turn, has been shown to contribute to depolarization (Beam, Hutchens, & Hmielowski, 2018).

While these examples provide initial evidence on ideological homogeneity in online networks (e.g., are Democrats more likely to be connected to Democrats?), a focus on ideology can only serve as a proxy for the extent to which individuals encounter views that are dissimilar to theirs. When it comes to analyzing the connection between similar and dissimilar stances, it seems more informative to focus on specific politically and civically relevant topics that are factually debated, that is, on the content of the discussion.

Against this backdrop, the present study proposes an analytical approach that addresses specific issue-related discussions on social media and opinion-based homogeneity therein. Accordingly, we refer to opinion-based homogeneity as the extent to which a set of political opinions that are similar are connected with each other (relative to the extent to

which they are connected to dissimilar opinions). While ideological homogeneity operates on a general group level in terms of being, for example, liberal or conservative, opinion-based homogeneity requires a reference to specific political topics. This topic-oriented approach is thought to offer a more nuanced view of the nature of homogeneous versus heterogeneous online discussions and the prevalence of like-minded spaces when it comes to political discussions.

To our knowledge, no research has addressed online homophily based on opinion-based homogeneity by combining natural language processing and social network analyses. Using the amalgamation of these two approaches, this study investigates to what extent citizens' opinion expressions in the form of user-generated comments are related to each other when they represent a similar stance on a politically relevant question. To this end, written German user-generated comments on political issues were analyzed.

Literature in this area has been limited to the investigation of social media platforms such as Facebook and Twitter, largely neglecting the most popular video-sharing platform YouTube. According to the website ranking platform SimilarWeb¹, YouTube is visited more often (28.9 billion visits in the last six months, as of November 2019) than Facebook (24.6 billion), and significantly more than Twitter (4.6) or Instagram (4.1). YouTube is turning more and more into a platform where users not only watch videos, but especially young users form communities to discuss videos or topics, and exchange opinions on current politically relevant debates (YouGov & BRAVO, 2017). Thus, it seems a pressing need to investigate the potential existence of political like-mindedness on the social platform YouTube. To formalize the general objectives of this paper, two questions guide this research:

RQ1. How high is the prevalence of opinion-based homogeneity among YouTube comments on specific political topics?

¹ <https://www.similarweb.com>

When addressing online homogeneity, there might still be differences between homogeneity at a large scale, referring to the whole network (e.g., the whole platform) which covers the full range of the topical discussions, and sub-networks in which discussions are based on reciprocal responses. Consequently, we ask:

RQ2. How does opinion-based homogeneity vary between analyses on a macro level (i.e., focusing on discussions across the full network) and a micro level (i.e., focusing on sub-networks) among YouTube comments?

To address these questions, this paper presents a combined approach of social network analysis (SNA) and sentiment analysis (SA). Crawling a multi-content social networking platform such as YouTube allows us to create a model based on unstructured German YouTube comments to predict the sentiment score of multiple users toward specific controversial topics. In particular, the present approach uses support vector machines (SVM) to predict the sentiment score on German comments of controversial political discussions on YouTube. These analyses were run for three different politically relevant topics: the right of same-sex couples to adopt children, a ban on headscarves, and climate change. These topics have been discussed extensively in the public and represent good examples of divisive issues that are associated with fundamental moral questions.

Background

Political Homogeneity in Online Communication

In many instances, it has been suggested that politically and civically relevant communication on social media can hold individual users captive in spaces in which they are exposed to political views that are in line with their pre-existing opinions (i.e., so-called “echo chambers”) (Boutyline & Willer, 2017; Sunstein, 2017). In light of democratic ideals, politically homogeneous spaces are assumed to lead to political polarization and radicalization since users are allegedly caught in self-reinforcing networks which, in the long run, could become more extreme (Prior, 2007). When it comes to analyzing whether and how individuals might get “caught” in those like-minded networks, different (non-mutually exclusive)

scenarios are conceivable (Flaxman, Goel, & Rao, 2016; Geschke, Lorenz, & Holtz, 2019):

(a) users actively homogenize their network and, therefore, their information sources, (b) algorithms shape the ideological environment of users, or (c) users are incidentally exposed to a thread of like-minded information (e.g., when comments refer to other comments that are uniform in the stance they express). The present work focuses on the latter scenario and investigates to what extent user-generated comments on political questions are related to congenial comments by others.

Initial evidence focusing on political homogeneity online showed that people are indeed connected to like-minded users to a larger extent than to politically opposing users in the United States (e.g., Bond & Messing, 2015; Boutyline & Willer, 2017). Theoretically, this pattern can be explained by the notion of selective exposure (Colleoni, Rozza, & Arvidsson, 2014; Knobloch-Westerwick, 2014; Zillmann & Bryant, 1985): People experience positive emotions when consuming information that conforms to their pre-existing views and feel stressed when the information contradicts their views. As a result, they seek out situations in which they are exposed to information that is in line with their views. This makes them more likely to affiliate with like-minded others and create homogeneous groups. While social media users may commonly be fully in control of their virtual acquaintances (e.g., in terms of friending or following someone or a news channel), they may not have full control over the information and stances they are exposed to incidentally, for instance, when browsing through certain Facebook or YouTube news channels (Lu & Lee, 2018). Following this logic, it seems worthwhile to ask to what extent users are actually exposed to and in contact with opinions they disagree with.

Empirical research addressing the potential existence of echo chambers in online networks has been based on two different approaches: On the one hand, survey research has relied on subjective estimates by social media users. This line of research, asking participants how frequently they are exposed to opinion or ideological diversity, has shown that on social media, people are incidentally exposed to heterogeneous opinions (e.g., Kim, 2018; Lee,

Choi, Kim, & Kim, 2014; Lu & Lee, 2018; Vaccari et al., 2016). On the other hand, another series of studies used observational data and made use of computational methods, especially focusing on SNAs.

The Assessment of Political Homogeneity Online based on Social Network Analyses

As a widely used method in research focusing on political homogeneity, SNA examines the properties of social networks – networks composed of people and their social connections with one another. In SNA, the property of an individual to seek social connections to other individuals with similar characteristics is called homophily. In other words, homophily is described as “the principle that a contact between similar people occurs at a higher rate than among dissimilar people” (McPherson, Smith-Lovin, & Cook, 2001, p. 416).

Several studies have therefore used network analysis to examine political like-mindedness in network data. Bakshy et al. (2015) analyzed Facebook data to examine the political homogeneity of friend networks to identify whether users read and share messages that are more consistent with their political ideological beliefs than cross-cutting content. Their findings showed that about 20% of users' Facebook friends were from the opposing party, which increases the probability that users will receive content that diverges from their own ideology. Another study focused on Twitter data to determine ideological homogeneity by analyzing 3.8 million Twitter users and a dataset of almost 150 million tweets on political and non-political topics (Barberá, Jost, Nagler, Tucker, & Bonneau, 2015). Their results revealed that Democrats were significantly more likely than Republicans to be involved in the cross-ideological dissemination of political and non-political information. Recently, Del Valle and Bravo (2018) ran a SNA of the Twitter network among Catalan parliamentarians and how information flows among them. Their study found that representatives are more likely to interact with members of their own party who share the same political interests.

On a methodical level, to identify echo chambers which are characterized by “disproportionate connections among ideologically similar political communicators” (Jasny,

Waggle, & Fisher, 2015), the structural properties of the social network – specifically, the likelihood of connections between members of a group – need to be compared with the political views of the members of the network. If the two are related, that is, if there is a group of individuals with a disproportionately high density of intra-group connections compared to the number of outside connections, whose members share political views that they do not share with non-members, this group can be considered an echo chamber, that is, a politically homogeneous communication space.

The identification of political homogeneity, thus, requires two steps: identifying a group (i.e., a subset) of users who agree politically, and measuring whether there is a disproportionately high number of connections between group members. Researchers have used various methods for both steps.

A useful way of quantifying the relationship between intra-group and inter-group connections is the E-I index. It was presented for the first time by Krackhardt and Stern (1988) and compares the strength of internal connections between members of a class to the strength of external connections to non-members. Other studies examining political echo chambers used similar methods and based their conclusions on the E-I index, e.g., to assess the fragmentation between pairs of discussion networks or the effect of tie strength on the polarization in such networks (Bright, 2018; Chan & Fu, 2017). By using the E-I index, it is possible to quantify the degree to which members of a group interact with each other, as opposed to interacting with others outside the group.

In addition, the identification of political homogeneity requires information about the political affiliations or views of the members of a social network. With observational network data at hand, the arguably most accurate source for inferring an actor's political views is the set of posts and comments in which he/she expressed his/her viewpoints. Working with unstructured text data poses especially difficult challenges (Stieglitz, Mirbabaie, Ross, & Neuberger, 2018), but there are a few studies that have used methods from natural language processing to tackle this problem.

Identifying Political Opinions based on Sentiment Analyses

Natural Language Processing (NLP) is a branch of Artificial Intelligence (AI) that deals with the interaction between human language and computers to allow them to understand incoming information and process it independently. One subdivision of NLP is called Sentiment Analysis (SA), also known as opinion mining. Machine learning approaches to SA classify texts by identifying their sentiment based on previously learned patterns. Machine learning “addresses the question of how to build computers that improve automatically through experience” (Jordan & Mitchell, 2015, p. 255). SA is a common tool to summarize emotional communication patterns on social media and is becoming increasingly important in the field of social media analytics (Stieglitz et al., 2018). Sentiment analyses are ideally suited to address the distribution of positive and negative viewpoints on a question of interest.

This method is of particular interest for the identification and further investigation of political homogeneity online. With this approach, it is possible to recognize whether and which people expressed a positive or negative stance on an issue and whether users are referring to each other. There are only few studies on political homogeneity which use machine learning approaches to infer the political views of users from the content of their messages. Colleoni et al. (2014) classified Twitter users as either political or nonpolitical (based on training data from blog posts) and as either Democrat or Republican (based on training data from users’ tweets). Their results suggest that the degree of homophily varies by political orientation: Democrats were less likely to have outbound ties to Republicans than Republicans to Democrats. Studies employing similar methodical approaches found that users are more likely interact with those who express similar views or stances than with those voicing dissimilar views (Himmelboim et al., 2016; Williams, McMurray, Kurz, & Lambert, 2015).

While previous research, therefore, offers initial evidence on how expressed sentiments are spread all over a network, most previous studies investigated the Twitter

network in which users are explicitly connected to each other (by the feature of “following”) and this original connection might be subjected to selective exposure tendencies (i.e., getting virtually acquainted only to those who are politically similar). Still, it has been left open how users respond to each other on particular issues on platforms that have less structured networks (e.g., YouTube), increasing the chance of getting exposed to counter-attitudinal content. For this purpose, it is necessary to a) focus explicitly on discussions about specific political topics and b) analyze the network and its sub-networks that are formed based on these topical interactions.

The Present Approach: A Combination of Sentiment and Social Network Analyses to Assess Opinion-Based Homogeneity

So far, we are unaware of approaches in which homogeneity is applied to individual topics and simultaneously combined with automated content analysis and SNAs. Previously, the determination of homogeneity was based on the basis of ideological classifications (i.e., the network patterns among liberals versus conservatives). We are only aware of few studies which examined polarization on the basis of topic-oriented approaches (e.g. Chan & Fu, 2017; Häussler, 2018). However, as public opinion forms based on issue-related discussions, it is key to focus on the analysis on specific topics. To this end, a combination of NLP – more precisely, automated SA – and SNAs is necessary. The present approach is structured as follows: First, based on manually labeled comments, the SA is performed with an SVM to predict the opinion climate for the entire network. Second, the results of the SA are then transformed into a network structure to compute the opinion-based homogeneity using the E-I index.

Method

Dataset

All data in this study were collected using a custom developed Python application which is directly connected to the YouTube API. Our application is able to collect multiple datasets by requesting queries from the internal YouTube search list, the video list, the

comments list, and the replies list of each individual video. Each request to the respective list has its own URL that allows the API to be accessed and data to be collected². For each list, we stored the requested data in a relational database.

The collected data contain the comments and replies of three controversial topics in Germany: “Kopftuchverbot in Deutschland” (headscarf ban in Germany), “Adoption für homosexuelle Paare” (adoption for same-sex couples) and “Klimawandel” (climate change) which also served as search queries. All of these topics are associated with political questions on which members of society have offered different answers. It has been suggested that especially morally loaded and controversial topics imply the potential to elicit processes of homogenization of opinion climates over time (Noelle-Neumann & Petersen, 2004). Accordingly, we believe that opinion-based homogeneity is more likely to be prevalent when focusing on such political topics (see Appendix³ A for more information about these topics).

When requesting the videos via search list, the parameter “relevantLanguage” was set to the value “de” in order to get primarily German content. Furthermore, we sorted the search queries for videos according to their relevance using the parameter “order,” whereas the parameter value is set to “relevance.” While the two datasets “adoption rights” and “headscarf ban” were acquired on May 15, 2018, the dataset on “climate change” was collected on January 2, 2019. Each dataset contains the user-generated comments as well as associated replies.

Table 1 provides an overview of the crawled videos with their corresponding search term and the data provided by this crawling. To analyze a more accurate selection of videos that reflect political issues, we filtered the videos by a specific categoryID⁴. In this case, a categoryID of 25 indicates the category of “Politics and News” in the YouTube API.

² Example URL for a search query on the keyword climate change:
<https://www.googleapis.com/youtube/v3/search?part=snippet&relevantLanguage=de&order=relevance&maxResults=50&climate+change&key=API-KEY>

³ https://osf.io/e92n3/?view_only=95ece274e9b74cc29dcadb49a06062fb

⁴ <https://developers.google.com/youtube/v3/docs/videoCategories>

Table 1. *Crawled YouTube videos.*

Search keyword	total results	total likes	total dislikes	total views	total comments	filtered comments
Adoption for same-sex couples	266	31,876	8,509	2,576,318	15,889	8,443
Headscarf ban in Germany	320	199,912	26,393	7,247,958	48,354	14,277
Climate change	336	167,236	16,136	10,387,029	46,894	18,185

Classification of opinions in social media

Manual Labeling. We created a human-annotated gold standard to create a sample of the 4,000 German YouTube comments for each topic by defining a coding scheme. This scheme ensures that the unlabeled data can be assigned to a unique class which represents the sentiment of the message. We use the term “sentiment” referring to comments expressing a positive or negative stance towards a specific topic (e.g., if a comment states “I hate headscarves,” this comment is classified as having a “negative” opinion of headscarves). This does not apply to comments whose general tone is positive or negative if they do not explicitly express a stance on the respective controversy.

We selected two well-trained independent annotators who received the same dataset with 4,000 randomly selected comments and replies for each topic. The data were labeled considering three mutually exclusive classes: negative, positive, and others. The coding scheme with corresponding topics and the listed classes is represented in Appendix A.

Agreement between the two annotators was measured using Krippendorff’s alpha (Hayes & Krippendorff, 2007). The value of 0.63 was obtained for 3-class annotation of the adoption rights data, whereas a value of 0.67 was obtained for the headscarf ban data. In the case of the climate change dataset, a value of 0.54 was determined. All inter-annotator agreement values are valid for further processing. To ensure better results for the machine learning model, we decided to use only those comments for further analysis on which both annotators agreed. This strategy guarantees that the sentiment can be clearly assigned to a

unique class without inconsistencies. Table 2 shows the distribution of sentiment classes for each of the three datasets.

Table 2. Labeled datasets indicating the distribution of different classes.

Sentiment	Dataset		
	Adoption rights	Headscarf ban	Climate change
Negative	339	400	416
Positive	530	294	356
Others	2432	2769	2328

To derive the impact of the data showing a disagreement between the annotators (borderline cases), we later projected these data onto our trained model to determine to what degree our model takes this borderline case into account. The results of the contingency table and the graph can be found in Appendix B.

Data Pre-Processing. We implemented multiple data pre-processing steps that structure and clean the data to decrease the level of noise in the subsequent analyses. These steps were the creation of a training (80%) and a testing set (20%), their cross-validation, and the transformation of cleaned comments to Term Frequency-Inverse Document Frequency (TF-IDF) vectors (for more information about the data pre-processing see Appendix C).

Support Vector Machine (SVM). The application of SVM in text classification or SA has been successfully carried out in many studies. A recent study used the in-memory framework Apache Spark to apply a SA by using an SVM with a rbf kernel to classify microblog comments (Yan, Yang, Ren, Tan, & Liu, 2017). Al-Smadi, Qawasmeh, Al-Ayyoub, Jararweh, and Gupta (2018) compared the performance of recurrent neural networks (RNN) and SVMs on a comprehensive aspect-based SA of Arabic hotel ratings. The results indicate that the SVM performs superior to the deep RNN in terms of the research tasks (aspect category identification, aspect opinion target expression, and aspect sentiment polarity identification). However, the use of SVM combined with the network method to measure homogeneity/heterogeneity in online networks is novel. The results of the above-mentioned

studies were very promising, and the performance of the classifiers was very high. Therefore, we decided to adapt them as a basis for our research.

The training of the SVM is realized by a pipeline (fixed sequence of steps) which starts by importing the cleaned training dataset and transform the text data into numerical feature vectors to make them readable for the algorithm. We used a bag-of-words approach of assigning each word to an integer and returning a vocabulary dictionary in the form of a document-term matrix. The pipeline ends by fitting the TF-IDF vectors in the SVM. Combining the processes of 5-fold cross-validation and grid search, we can initialize different parameters during training and localize the best combination of parameters for each fold separately. The best parameter set is used which reaches the highest subjective F1-score.

The F1-score is used to determine the performance of the model. Especially when the class distribution is uneven, it is more precise than the simple accuracy measure. In a systematic test, we used an SVM with a linear kernel based on the LIBSVM implementation (Chang & Lin, 2011) of scikit-learn (Pedregosa et al., 2011). The optimization of the parameters was carried out through a grid search in 5-fold cross-validation. Instead of only tuning the parameters of the classifier, we also tuned parameters that deal with the process of data pre-processing. The list of all tuned hyperparameters is given in Appendix D.

The evaluation of the final model with their optimal parameters is based on the unseen test dataset. We apply the weighted F1-score as the metric to measure the performance of the model. Table 3 reveals the results of the prediction on the test dataset with their metrics.

Table 3. Summary of the precision, recall, F1-score for each class.

Topic	Sentiment	Metrics			
		Precision	Recall	F1-Score	Support
Adoption rights	Negative	0.62	0.52	0.56	64
	Positive	0.61	0.75	0.67	108
	Others	0.93	0.91	0.92	489
	Weighted avg.	0.85	0.84	0.85	661
Headscarf ban	Negative	0.56	0.53	0.54	76
	Positive	0.72	0.54	0.62	57
	Others	0.93	0.96	0.95	560

	Weighted avg.	0.88	0.88	0.88	693
Climate change	Negative	0.67	0.63	0.65	99
	Positive	0.49	0.41	0.44	69
	Others	0.91	0.95	0.93	452
	Weighted avg.	0.83	0.84	0.83	620

The normalized confusion matrices for the three datasets (see Figure 1) shows that the F1-score is strongly driven by the “others” category. Precision and recall for the positive and negative categories are lower. In summary, the confusion matrices show that the biggest performance losses are due to the classes positive and negative. In both classes, data is likely to be classified in the opposite category which may be due to the low amount of training data.

Figure 1. Normalized confusion matrix across all classes.

Since the classifier achieves valid predictions on the test dataset and an adequate F1-score of 0.85 for the adoption dataset, 0.88 for the headscarf ban dataset, and 0.83 for the climate change dataset, it was to predict the sentiment of the comments across the whole dataset.

Some users wrote several comments, each of which may express a combination of stances. To simplify the visualization of the network structure and the calculation of homogeneity, each user was assigned exactly one class as follows. Platt scaling was used to generate probability estimates for each class and comment (Chang & Lin, 2011). In order to summarize these values, we calculated, for each user, the average probability of each class across their comments. The user was assigned to the class that was the most likely on average. For an overview of the distribution of the prediction, please see Appendix D.

Building a Network on YouTube

In contrast to other social networking sites such as Facebook and Twitter, in which friendship requests and follower relationships play an integral role, the structure of the social network of YouTube users is not nearly as visible. Users can interact by commenting on videos and by commenting on other users' comments. In this study, we examined the interactions between users' comments, associated replies, and users who uploaded the video. Thus, the focus lies on the exchange of messages between users. The SNA is structured in three parts. The first part is the creation of the network using the YouTube data to visualize the interactions across all videos (see Appendix F). Statistics are used to provide a general overview of the network and to detect any conspicuous features. The second part deals with the computation of opinion-based homogeneity with the Krackhardt E-I ratio of the global network. The last part of the network analysis includes the segmentation of the networks into smaller sub-networks using the fast-greedy algorithm and the calculation of opinion-based homogeneity on a macro level (covering every comment on YouTube on that topic) as well as exchanges on the micro level (in sub-networks).

As we aim to identify the extent to which users have varying opinions on a particular topic, we decided to exclude the category "others" from the analysis as well as the removal of self-links. Topic modeling was used to gain an overview of the data that was thus discarded (Appendix E). The results show that the comments in this category were off-topic and therefore do not directly contribute to the discussion between proponents and opponents on

the three controversies. The removal of these off-topic posts from the network led to the creation of isolated nodes that no longer had any connections to other nodes and, therefore, had a degree equal to zero. These nodes were also deleted from the network. To ensure that the results with the class “others” are not entirely ignored, we also performed the entire analysis on all three networks including this category. The findings can be found in Appendix F.

To understand and explore the network more closely and to gain a deeper insight, we have calculated various statistics and reported the results for the three datasets in Table 4.

Table 4. Network properties.

Network parameter	Datasets		
	Adoption rights	Headscarf ban	Climate change
Nodes	536	968	626
Edges	523	1064	703
Avg. degree	0.98	1.10	1.12
Diameter	3	3	4
Max out-degree	8	18	87
Max in-degree	469	615	300
Density	0.0018	0.0011	0.0018

The characteristics of all directed networks shown in Table 4 demonstrate that the average degree is about one, suggesting that a typical user interacts with approximately one other user. In general, accounts that have uploaded a video that many other users have commented on have a higher in-degree (comments addressed to them). In addition to the in-degree, the out-degree shows which users have interacted with other users the most frequently by writing a comment. The low density values might be explained by the fact that the data originates from a real network in which the users are not linked by friendships but by their comments to each other, as well as by the high number of nodes. This pattern seems plausible in a public network where the investigation and the focus is on comments. The combination of in-degree, out-degree as well as number of nodes and edges explain the difference in the diameter.

Measuring Opinion-Based Homogeneity. One of the main goals of this study is the measurement of opinion-based homogeneity based on the sentiment of comments. To measure the degree of homogeneity, the E-I index is an appropriate choice. The formula of the global E-I Index is defined as follows:

$$EI\ Index = \frac{E - I}{E + I}$$

where E is the number of external links to a given subgroup (sentiment) and I is the number of internal links to or between nodes within that subgroup (sentiment).

The index is in a range of -1.0 to +1.0. A value of -1.0 indicates that the network is entirely homophilous with respect to the classes, i.e., all connections in the network are between members of the same class (alternatively, each connected component in the graph only involves members of the same class.). A value of +1.0 indicates an entirely heterophilous network in which there are no connections between members of the same class (i.e., a multipartite graph). In addition to measuring the global homogeneity of the network, it is possible to compute a homogeneity value for each specific class (or sentiment) to identify which sentiment has characteristics of a homogeneous interaction cluster. For example, the E-I index of the negative class would be -1.0 if all connections, both incoming and outgoing, that involve a member of the negative class were links to members of the same class. It would be +1.0 if there were no direct connections between any two members of the negative class. The index has previously been used in studies to investigate homogeneity in offline networks (Eveland & Kleinman, 2013; Levendosky et al., 2004). To clarify the interpretation of the E-I index, Appendix F shows three networks with different properties.

Identification of Communities and Extraction of Sub-Networks. The detection of sub-networks to calculate the opinion-based homogeneity of each community could give further clues about the opinion climate and possible differences between the macro and the micro level. Especially in sociology, it is necessary for many activities to identify the internal structures and groups of social networks. However, this can also be applied to online social

media such as YouTube, Facebook, or Twitter in order to recognize the community structure of a network of users.

For this study, we used the fast-greedy algorithm introduced by Newman (2004) and Clauset, Newman, and Moore (2004) which is a hierarchical approach for the optimization of modularity in network analysis. This algorithm has already been applied to social network data from Twitter in several studies (e.g., Mercea & Yilmaz, 2018) and has also achieved the best results in the area of community detection based on modularity (Bello-Orgaz, Hernandez-Castro, & Camacho, 2017). The goal of this technique is to optimize the modularity to find community structures in the network. The higher the modularity score, the better is the sophisticated internal structure of the network represented. To determine the algorithm, we compared fast-greedy on a test basis with two other algorithms called Walktrap (Pons & Latapy, 2006) and Louvain (Blondel, Guillaume, Lambiotte, & Lefebvre, 2008); the results can be found in the Appendix G.

Results

Figures 2-4 show the graphical representations of all three topic networks. The nodes in the network represent individual users of YouTube, i.e., users who have written comments, users who have responded to comments, and channel owners, some of whom have also written comments or replies. The color of the nodes represents their sentiment score: red for negative, green for positive, and black for channel owners who have not written any comments and are only in the dataset because they uploaded a relevant video.

Figure 2. Discussion network on the topic of adoption rights for same-sex couples.

Figure 3. Discussion network on the topic of the headscarf ban.

Figure 4. Discussion network on the topic of climate change.

Due to the aggregated probability values of the individual classes, it is easy to detect which opinion the users represent. The connections of the individual nodes to each other reflect their interaction in the form of comments. It should be noted here that this is a directed network, so it is possible to see the direction of the information flow. The hubs in the network represent channel owners who uploaded the videos that many users commented on. Furthermore, it can be seen that apart from the hubs, the connections to the individual nodes are distributed in a very mixed way and, thus, a heterogeneous opinion climate prevails.

Table 5. Properties of opinion-based homogeneity.

		Network statistics			
	Sentiment	Internal Ties	External Ties	Class E-I Index	Global E-I Index
Adoption rights	Negative	31	173	0.70	0.72
	Positive	41	278	0.74	
Headscarf ban	Negative	194	621	0.52	0.58
	Positive	28	221	0.78	
Climate change	Negative	102	320	0.52	0.61
	Positive	34	247	0.76	

Looking at the classes for each topic, it is evident that YouTube users more often comment on messages that express an opinion that is different to their own than on messages with a similar stance. This is corroborated by the E-I index which approaches +1.0 and the relatively small number of internal ties (see Table 5). In addition to the visualization of the entire network, the three largest sub-networks are presented graphically in Figure 5. The visualization of the sub-networks gives a more detailed view of the network because it offers evidence about the opinion-based homogeneity related to videos with a higher number of comments.

Figure 5. Sb-networks.

(a) Adoption rights.

(b) Headscarf ban.

(c) Climate change.

By dividing the network into sub-networks, the individual communities can be examined more precisely, i.e., structures of single or several channel owners are recognized more effectively.

When comparing the three largest sub-networks of each dataset, it is noticeable that sub-networks on the topics “headscarf ban” and “climate change” have a higher number of users responding to comments. This is in line with the significantly higher number of comments related to those topics. Furthermore, both topics are marked by denser network structures in which different channel owners are linked by users.

The sub-communities are relatively large, and they do not reflect homogeneous opinion climates with users unanimously speaking out in favor of or against a political decision. Instead, they show a moderately diverse exchange of opinions. By examining the sub-networks, a significantly more precise analysis and results can be created for the micro-level where not specific sections of the whole network are visible. It can be noted that every sub-network exhibits heterogeneous behavior with regard to the opinion climate. Table 6 shows the results for sub-networks of global as well as class E-I Indexes.

Table 6. Properties of opinion-based homogeneity – Sub-networks.

Dataset	Sub-network	Sentiment	Statistics			
			Internal Ties	External Ties	Class E-I Index	Global E-I Index
Adoption rights	I	Negative	1	15	0.88	0.92
		Positive	2	58	0.93	
	II	Negative	9	19	0.36	0.61
		Positive	0	18	1	
	III	Negative	1	12	0.85	0.94
		Positive	0	22	1	
Headscarf ban	I	Negative	30	182	0.72	0.77
		Positive	0	49	1	
	II	Negative	28	71	0.43	0.59
		Positive	1	40	0.95	
	III	Negative	42	71	0.26	0.35
		Positive	0	17	1	
Climate change	I	Negative	2	48	0.92	0.89
		Positive	3	39	0.86	
	II	Negative	8	35	0.63	0.79
		Positive	1	41	0.95	
	III	Negative	4	26	0.73	0.61
		Positive	11	36	0.53	

Discussion

The present work was intended to (a) offer a new methodological approach to address opinion-based homogeneity using a combination of NLP and SNA and (b) provide preliminary evidence on the prevalence of opinion-based homogeneity regarding three (politically) controversial topics discussed on the platform YouTube.

Addressing RQ1, results based on the combination of NLP and SNA did not offer evidence for opinion-based homogeneity regarding positively and negatively valenced YouTube comments on the topics of adoption rights for same-sex couples, the prohibition of headscarves, or climate change. Instead, we found a moderate level of opinion-based heterogeneity when it came to the connection, that is, cross-references among user-generated comments on YouTube. In other words, comments on these three political issues were more likely to be connected to dissimilar than to similar comments.

Regarding RQ2, it can be concluded that there are only minor differences between the macro and the micro level in the determination of opinion-based homogeneity versus heterogeneity. Both analyses – either focusing on the whole network or on sub-network – show similar structures. In particular, a closer look at the different sub-networks can lead to a more precise analysis because structures of individual communities can be focused, and opinion-based homogeneity can be calculated specifically. Given analyses at both levels, one cannot conclude that users on YouTube are exposed to a series of connected messages that all represent like-mindedness in terms of a uniform opinion climate. This result challenges previous research offering evidence for the, albeit weak, prevalence of ideological homogeneity of social networks such as Twitter (Bakshy et al., 2015; Barberá et al., 2015; Boutyline & Willer, 2017; Colleoni et al., 2014). These studies, however, focused on ideological homogeneity, that is, to what extent Democrats and Republicans interact with each other on platforms such as Twitter. Political discussions, though, may become diverse and include diverging viewpoints even within these ideological clusters. Moreover, as indicated by the same line of research, users still have ties to “the other side.” While previous studies

assumed that due to their cross-ideological connections, social media users might encounter content that is created or published by an ideologically deviant source (Bakshy et al., 2015; Barberá et al., 2015), it remained unclear whether users indeed encounter cross-cutting content. The present work provides initial evidence that users' opinion expressions are more likely to be associated with divergent than with congenial comments by others. In fact, this pattern is in line with the notion of "corrective action" (Rojas, 2010) stating that users feel encouraged to become outspoken online when they feel that their opinion is underrepresented. According to the patterns found on YouTube, this seems to apply as users tend to voice their political stance especially in relation to previous comments that were different to their opinion.

The only group with significantly more in-group interactions than out-group interactions, as evidenced by a negative class E-I index, is the "others" group (see Appendix F). This group consists of users that discuss topics that are only vaguely related to the controversy in question (see Appendix E). From the results of the study, it appears that such comments commonly trigger a similarly off-topic response, leading to the creation of entire comment threads that diverge from the topic of the video. These groups are therefore homogeneous, but not with respect to their opinion on the topic of the video, which would be a prerequisite for the existence of opinion-based homogeneity in the sense of the present research questions.

The combination of machine learning and SNA allowed measuring opinion-based homogeneity by assigning opinions to a particular class, training a model based on these labelled data, and applying this model to all comments. Still, it should be noted that the values predicted by means of machine learning do not reach perfect accuracy due partly to the size of the dataset and the unequal number of samples for the different sentiment classes, especially for the over-represented class "others." However, the prediction of the test datasets gives us a rough impression of the extent to which the classification works well on previously unseen data and whether the model has generalized well or only classifies examples correctly that

closely resemble the training data. Looking at the performance metrics, it can be seen that the model generalizes well with class weights that are suitable for rebalanced datasets.

In general, unbalanced datasets are a common problem in machine learning contexts which can be solved by crawling and labeling even larger and more balanced datasets to improve data quality and provide more training data for the model. In the pre-labeling procedure, we have also helped to improve data quality by only using records for analysis where both annotations matched. It is remarkable that most comments crawled on all three topics did not elaborate on the question of interest. This is in line with early research evaluating the deliberative ideals of online discussions which assessed that many contributions made by users are off-topic (Janssen & Kies, 2005; Min, 2007; Schneider, 1996). Consequently, while the present findings may allow us to be optimistic about the heterogeneity of political discussions on YouTube, it raises concerns about the relative weight of these on-topic exchanges in face of a huge number of off-topic interactions.

Limitations

Our method of crawling YouTube comments about three different topics does not represent the full landscape of the political discussion on this platform but rather gives an overview of three currently discussed debates and exchanges to determine the degree of homogeneity. One reason for this is the limitation of the YouTube API which only enables crawling a fixed number of comments and videos.

Another limitation which can affect the opinion climate in the analysis, is the imbalance of the labeled classes, making the training more challenging. Using 5-fold cross validation and class weights which can be used for addressing the generalization problem as well as for the hyperparameter search, we have tried to prevent the model from overfitting. However, one reason why the accuracy of this model is so high is that this over-represented class is more common in the training and test datasets, and it is therefore also predicted more often automatically. This also means that the accuracy of the individual models strongly depends on the available data. This, in turn, has a direct influence on the calculation of

opinion-based homogeneity in the network. Increasing the amount of data would therefore also lead to the creation of a separate validation dataset which in the analysis could reduce further knowledge about the prediction and the accuracy of the model as well as the reduction of the overfitting.

As a further limitation of the work, it should be mentioned that excluding the cases of disagreement between both annotators can influence the result of the classification and give a misleading impression of the accuracy of the classifier. To prevent this, a higher number of annotators would be necessary in order to have a uniform understanding of the comments and therefore increase the precision of the trained model. The results of the contingency table in Appendix B show that most of these borderline cases belonged to the class “others,” which is also the most frequently represented one in the dataset.

For the present study, the YouTube network was built based on the connection of videos, comments, and replies. Consequently, the network does not show the full connection structure between the individual users (e.g., friendships on Facebook). Accordingly, we consider homogeneity in the discursive sense between users although the criteria according to which the user selects individual videos or channels cannot be determined on the basis of this structure. In the present work, video uploaders assumed a key role as their opinion (provided that they expressed one) was a central connection node in the networks. Their stance was inferred from any comments they had made on their own and others’ videos. Future research could also take the role of the video itself and its stance on the political question into account and investigate its interplay with the opinion climate that emerges in the related comment section.

While this applied approach has been limited to the YouTube platform, it is possible to apply the same approach to other social networking platforms such as Facebook or Twitter (using further political topics, in other language contexts) to measure opinion-based homogeneity there as well. A systematic comparison of homogeneity across different social media services will contribute to developing a robust understanding of the dynamics of

political discussions online and the factors that determine whether they become homogeneous or heterogeneous.

Conclusion & Further Work

This study has developed an approach to measure opinion-based homogeneity based on textual messages with SA and SNA techniques on the YouTube platform by evaluating three relevant and politically controversial topics. Specifically, we investigated, based on communication data on YouTube, how expressed opinions in the form of user-generated comments are connected to each other and to what extent opinion-based homogeneity and heterogeneity mark the political discourse. In contrast to ideological homophily, which is more suitable for the recognition of moral values and political identities, the present approach allows the investigation of dynamic opinion climates which can change in the course of political discourses.

The combination of the two methods SNA and SA has shown that a measurement of opinion-based homogeneity based on YouTube comments is possible and can also be adapted to different topical contexts and a variety of social platforms. In the overall network, instead of finding evidence for opinion-based homogeneity, we found a moderate level of connectivity among dissimilar opinions expressed in user-generated comments. Thus, comments who expressed either a positive or a negative stance toward one of the three political issues were more likely to be associated with a heterogeneous than with a homogeneous environment. A similar pattern was found when the whole network was divided into sub-networks, e.g., in which a lot of comments were related to each other. Accordingly, this paper contributes to computational communication research in three respects:

1. It offers a blueprint for a combination of computational methods (SA and SNA) that enable the analysis of large communication datasets in light of potential social dynamics (such as communication content becoming homogeneous).

2. While previous network analyses focused predominantly on Twitter, this work relies on political communication content available on the platform YouTube, a platform that is growing as a political arena, especially for younger users.
3. Given the public debate about so-called echo chambers and political homogeneity in social media, this paper offers evidence based on automated analyses of observational data that extends previous research by not focusing on ideological homogeneity but on opinion-based and issue-related homogeneity.

Acknowledgments

This research was supported by the Digital Society research program funded by the Ministry of Culture and Science of the German State of North Rhine-Westphalia (Grant Number: 005-1709-0004), Junior Research Group “Digital Citizenship in Network Technologies” (Project Number: 1706dgn009).

References

- Al-Smadi, M., Qawasmeh, O., Al-Ayyoub, M., Jararweh, Y., & Gupta, B. (2018). Deep recurrent neural network vs. support vector machine for aspect-based sentiment analysis of Arabic hotels' reviews. *Journal of Computational Science*, 27, 386–393. <https://doi.org/10.1016/j.jocs.2017.11.006>
- Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239), 1130–1132. <https://doi.org/10.1126/science.aaa1160>
- Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science*, 26(10), 1531–1542. <https://doi.org/10.1177/0956797615594620>
- Beam, M. A., Hutchens, M. J., & Hmielowski, J. D. (2018). Facebook news and (de)polarization: Reinforcing spirals in the 2016 US election. *Information, Communication & Society*, 21(7), 940–958. <https://doi.org/10.1080/1369118X.2018.1444783>
- Bello-Orgaz, G., Hernandez-Castro, J., & Camacho, D. (2017). Detecting discussion communities on vaccination in twitter. *Future Generation Computer Systems*, 66, 125–136. <https://doi.org/10.1016/j.future.2016.06.032>
- Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10), P10008. <https://doi.org/10.1088/1742-5468/2008/10/p10008>
- Bond, R., & Messing, S. (2015). Quantifying social media's political space: Estimating ideology from publicly revealed preferences on Facebook. *American Political Science Review*, 109(1), 62–78. <https://doi.org/10.1017/S0003055414000525>
- Boutyline, A., & Willer, R. (2017). The social structure of political echo chambers: Variation in ideological homophily in online networks. *Political Psychology*, 38(3), 551–569. <https://doi.org/10.1111/pops.12337>

- Bright, J. (2018). Explaining the emergence of political fragmentation on social media: The role of ideology and extremism. *Journal of Computer-Mediated Communication*, 23(1), 17–33. <https://doi.org/10.1093/jcmc/zmx002>
- Chan, C., & Fu, K. (2017). The relationship between cyberbalkanization and opinion polarization: Time-series analysis on Facebook pages and opinion polls during the Hong Kong Occupy Movement and the associated debate on political reform. *Journal of Computer-Mediated Communication*, 22(5), 266–283. <https://doi.org/10.1111/jcc4.12192>
- Chang, C.-C., & Lin, C.-J. (2011). LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3), 27:1–27:27. <https://doi.org/10.1145/1961189.1961199>
- Clauset, A., Newman, M. E. J., & Moore, C. (2004). Finding community structure in very large networks. *Phys. Rev. E*, 70(6), 066111. <https://doi.org/10.1103/PhysRevE.70.066111>
- Colleoni, E., Rozza, A., & Arvidsson, A. (2014). Echo chamber or public sphere? Predicting political orientation and measuring political homophily in Twitter using big data. *Journal of Communication*, 64(2), 317–332. <https://doi.org/10.1111/jcom.12084>
- Del Valle, M. E., & Bravo, R. B. (2018). Echo Chambers in parliamentary Twitter networks: The Catalan case. *International Journal of Communication*, 12, 21.
- Eveland, W. P., & Kleinman, S. B. (2013). Comparing general and political discussion networks within voluntary organizations using social network analysis. *Political Behavior*, 35(1), 65–87. <https://doi.org/10.1007/s11109-011-9187-4>
- Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, 80(S1), 298–320. <https://doi.org/10.1093/poq/nfw006>
- Geschke, D., Lorenz, J., & Holtz, P. (2019). The triple-filter bubble: Using agent-based modelling to test a meta-theoretical framework for the emergence of filter bubbles and

- echo chambers. *British Journal of Social Psychology*, 58(1), 129–149.
<https://doi.org/10.1111/bjso.12286>
- Graham, T. (2015). Everyday political talk in the internet-based public sphere. In S. Coleman & D. Freelon (Eds), *Handbook of digital politics* (pp. 247–263). Cheltenham, UK: Edward Elgar Publishing.
- Guo, L., Rohde, J. A., & Wu, H. D. (2018). Who is responsible for Twitter’s echo chamber problem? Evidence from 2016 US election networks. *Information, Communication & Society*. <https://doi.org/10.1080/1369118X.2018.1499793>
- Häussler, T. (2018). Heating up the debate? Measuring fragmentation and polarisation in a German climate change hyperlink network. *Social Networks*, 54, 303–313.
<https://doi.org/10.1016/j.socnet.2017.10.002>
- Hayes, A. F., & Krippendorff, K. (2007). Answering the call for a standard reliability measure for coding data. *Communication Methods and Measures*, 1(1), 77–89.
<https://doi.org/10.1080/19312450709336664>
- Himmelboim, I., Sweetser, K. D., Tinkham, S. F., Cameron, K., Danelo, M., & West, K. (2016). Valence-based homophily on Twitter: Network analysis of emotions and political talk in the 2012 presidential election. *New Media & Society*, 18(7), 1382–1400.
<https://doi.org/10.1177/1461444814555096>
- Janssen, D., & Kies, R. (2005). Online forums and deliberative democracy. *Acta Politica*, 40(3), 317–335. <https://doi.org/10.1057/palgrave.ap.5500115>
- Jasny, L., Waggle, J., & Fisher, D. R. (2015). An empirical examination of echo chambers in US climate policy networks. *Nature Climate Change*, 5(8), 782.
<https://doi.org/10.1038/nclimate2666>
- Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>
- Kim, M. (2018). How does Facebook news use lead to actions in South Korea? The role of Facebook discussion network heterogeneity, political interest, and conflict avoidance

- in predicting political participation. *Telematics and Informatics*, 35(5), 1373–1381.
<https://doi.org/10.1016/j.tele.2018.03.007>
- Knobloch-Westerwick, S. (2014). *Choice and preference in media use: Advances in selective exposure theory and research*. Routledge.
- Krackhardt, D., & Stern, R. N. (1988). Informal Networks and Organizational Crises: An Experimental Simulation. *Social Psychology Quarterly*, 51(2), 123–140.
<https://doi.org/10.2307/2786835>
- Lee, J. K., Choi, J., Kim, C., & Kim, Y. (2014). Social media, network heterogeneity, and opinion polarization. *Journal of Communication*, 64(4), 702–722.
<https://doi.org/10.1111/jcom.12077>
- Levendosky, A. A., Bogat, G. A., Theran, S. A., Trotter, J. S., Eye, A. von, & Davidson, W. S. (2004). The social networks of women experiencing domestic violence. *American Journal of Community Psychology*, 34(1–2), 95–109.
<https://doi.org/10.1023/B:AJCP.0000040149.58847.10>
- Lu, Y., & Lee, J. K. (2018). Stumbling upon the other side: Incidental learning of counter-attitudinal political information on Facebook. *New Media & Society*, 1461444818793421. <https://doi.org/10.1177/1461444818793421>
- McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1), 415–444.
<https://doi.org/10.1146/annurev.soc.27.1.415>
- Mercea, D., & Yilmaz, K. E. (2018). Movement social learning on Twitter: The case of the People’s Assembly. *The Sociological Review*, 66(1), 20–40.
<https://doi.org/10.1177/0038026117710536>
- Min, S.-J. (2007). Online vs. Face-to-face deliberation: Effects on civic engagement. *Journal of Computer-Mediated Communication*, 12(4), 1369–1387.
<https://doi.org/10.1111/j.1083-6101.2007.00377.x>

- Neubaum, G., & Krämer, N. C. (2017). Monitoring the opinion of the crowd: Psychological mechanisms underlying public opinion perceptions on social media. *Media Psychology*, 20(3), 502–531.
- Newman, M. E. J. (2004). Fast algorithm for detecting community structure in networks. *Phys. Rev. E*, 69(6), 066133. <https://doi.org/10.1103/PhysRevE.69.066133>
- Noelle-Neumann, E., & Petersen, T. (2004). The spiral of silence and the social nature of man. In *Handbook of political communication research* (pp. 357–374). Routledge.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Duchesnay, E. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Pons, P., & Latapy, M. (2006). Computing communities in large networks using random walks. *J. Graph Algorithms Appl.*, 10(2), 191–218.
https://doi.org/10.1007/11569596_31
- Prior, M. (2007). *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections*. Cambridge, NY: Cambridge University Press.
- Rojas, H. (2010). “Corrective” actions in the public sphere: How perceptions of media and media effects shape political behaviors. *International Journal of Public Opinion Research*, 22(3), 343–363. <https://doi.org/10.1093/ijpor/edq018>
- Schneider, S. M. (1996). Creating a democratic public sphere through political discussion: A case study of abortion conversation on the Internet. *Social Science Computer Review*, 14(4), 373–393. <https://doi.org/10.1177/089443939601400401>
- Stieglitz, S., Mirbabaie, M., Ross, B., & Neuberger, C. (2018). Social media analytics—Challenges in topic discovery, data collection, and data preparation. *International Journal of Information Management*, 39, 156–168.
- Sunstein, C. R. (2017). *# Republic: Divided democracy in the age of social media*. Princeton University Press.

- Vaccari, C., Valeriani, A., Barberá, P., Jost, J. T., Nagler, J., & Tucker, J. A. (2016). Of echo chambers and contrarian clubs: Exposure to political disagreement among German and Italian users of Twitter. *Social Media + Society*, 2(3), 1–24.
<https://doi.org/10.1177/2056305116664221>
- Williams, H. T., McMurray, J. R., Kurz, T., & Lambert, F. H. (2015). Network analysis reveals open forums and echo chambers in social media discussions of climate change. *Global Environmental Change*, 32, 126–138. <https://doi.org/10.1016/j.gloenvcha.2015.03.006>
- Yan, B., Yang, Z., Ren, Y., Tan, X., & Liu, E. (2017). Microblog sentiment classification Using parallel SVM in Apache Spark. *2017 IEEE International Congress on Big Data (BigData Congress)*, 282–288. <https://doi.org/10.1109/BigDataCongress.2017.43>
- YouGov, & BRAVO. (2017). *Politische Jugendstudie [Political youth study]*. Retrieved from https://campaign.yougov.com/DE_2017_07_Political_Bravo_Jugendstudie_DE_2017_06_Political_Die_Deutschen_und_die_Politik_Landing.html
- Zillmann, D., & Bryant, J. (Eds.). (1985). *Selective exposure to communication*. Hillsdale, N.J: L. Erlbaum Associates.