

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/344294839>

Impact and dynamics of hate and counter speech online

Preprint · September 2020

CITATIONS

0

READS

172

5 authors, including:



Keyan Ghazi-Zahedi

Max Planck Institute for Mathematics in the Sciences

45 PUBLICATIONS 433 CITATIONS

SEE PROFILE



Jean-Gabriel Young

University of Vermont

38 PUBLICATIONS 368 CITATIONS

SEE PROFILE



Laurent Hébert-Dufresne

Santa Fe Institute

108 PUBLICATIONS 1,108 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Quantifying Morphological Computation [View project](#)



POMDPs [View project](#)

IMPACT AND DYNAMICS OF HATE AND COUNTER SPEECH ONLINE

Joshua Garland*
Santa Fe Institute
1399 Hyde Park Road,
Santa Fe, NM 87501 USA
joshua@santafe.edu

Keyan Ghazi-Zahedi*
Max Planck Institute for Mathematics in the Sciences,
Inselstrasse 22, 04103
Leipzig, Germany

Jean-Gabriel Young
Department of Computer Science,
Vermont Complex Systems Center,
University of Vermont, Burlington, VT 05405, USA

Laurent Hébert-Dufresne
Department of Computer Science,
Vermont Complex Systems Center,
University of Vermont, Burlington, VT 05405, USA

Mirta Galesic
Santa Fe Institute
1399 Hyde Park Road,
Santa Fe, NM 87501 USA
galesic@santafe.edu

September 21, 2020

ABSTRACT

Citizen-generated counter speech is a promising way to fight hate speech and promote peaceful, non-polarized discourse. However, there is a lack of large-scale longitudinal studies of its effectiveness for reducing hate speech. We investigate the effectiveness of counter speech using several different macro- and micro-level measures of over 180,000 political conversations that took place on German Twitter over four years. We report on the dynamic interactions of hate and counter speech over time and provide insights into whether, as in ‘classic’ bullying situations, organized efforts are more effective than independent individuals in steering online discourse. Taken together, our results build a multifaceted picture of the dynamics of hate and counter speech online. They suggest that organized hate speech produced changes in the public discourse. Counter speech, especially when organized, could help in curbing hate speech in online discussions.

1 Introduction

Hate speech is rampant on many online platforms and manifests in many different forms, e.g., insulting or intimidating, encouraging exclusion, segregation, and calls for violence, as well as spreading harmful stereotypes and misinformation about a group or individuals based on their race, ethnicity, gender, creed, religion, or political beliefs [1–10]. While it is widely accepted that hate speech is a growing problem on online platforms, what to do about it is a point of contention. One proposal is to more or less automatically detect and remove hateful content. This is problematic for several reasons, including the nuanced and constantly evolving nature of hate speech, societal and legal norms about free speech, and the possibility of merely moving hate to other platforms rather than eliminating it [11]. An emerging alternative is citizen-driven *counter speech*, whereby users of online platforms themselves generate responses to hateful content in order to stop it, reduce its consequences, discourage it [12–14], and ultimately increase civility and deliberation quality of online discussions [15, 16]. Like hate, counter speech can take many forms, including providing facts, pointing

*Denotes equal contribution.




Figure 1: Examples of reply trees. Nodes are tweets and edges denote a “replied to” relationship. Here the node colors denotes whether the content of that tweet was hate (red), counter (blue) or neutral (white), with intensity of colors representing the magnitude of the hate score, which is defined in Section 4.2.

to logical inconsistencies in hateful messages, attacking the perpetrators, supporting the victims, spreading neutral messages, or flooding a discussion with unrelated content [17–23].

The effectiveness of counter-speech in curbing the spread of hatred online is not well understood. Until recently, there has been a lack of longitudinal, large-scale studies of counter speech [3, 24]. One reason has been the difficulty of designing automated algorithms for discovering counter speech in large online corpora. Past studies have provided insightful analyses of the effectiveness of counter speech, but were limited in scope as they relied on relatively small, hand-coded examples of discourse [15, 25–28]. Recently, however, we made use of a unique situation in Germany where two self-labeling groups engaged in organized online hate and counter speech to develop an automated classifier for labeling hate and counter speech in a very large corpora of sample text [29]—orders of magnitude bigger than any other study to date. Here we use this classification system to perform a longitudinal study of more than 180,000 fully-resolved Twitter conversations that took place from 2015 to 2018 within German political discourse to study the *dynamics* of hate and counter speech and gain insight into the potential effectiveness of counter speech.

We aim to gain insight into the effectiveness of counter speech by measuring how hate and counter speech interact with each other on macro and micro levels. On the *macro-level*, we study changes in the composition of online discourse, reflected in the prevalence of hate, counter, and neutral speech over time. This analysis helps assess the overall deliberation value and civility of discussions [15, 27, 30]. On the *micro-level*, we study interactions between hate and counter speakers, including how they support and reply to each other and how each particular utterance changes the overall subsequent discourse [26].

Drawing from the literature on bullying, which partially overlaps with hate speech [2], we anticipate that both organized counter and hate speech are more impactful than independent individual efforts. Indeed, the presence of peers and support are important for both bullying and countering behaviors. Bullies often seek an approving crowd when bullying, because attention justifies their behavior, helps stifle resistance, and enables them to achieve visibility and social status [31]. Bystanders, as well, often look to others when deciding whether to actively oppose it and help the victim, because others helping signals that the prevailing social norm opposes bullies, and because they can provide support and protection against bullies [32, 33]. Online hate and counter speech is similar to real-world bullying in that individuals who engage in either kind of speech can become targets of online hate themselves, and may even be threatened by physical violence. In addition, seeing that one is a lone countering voice in a flood of hateful or counter messages can be a powerful signal that the prevailing social norm opposes one’s own views and that the effort put in exposing one’s own view might be futile [34, 35]. Therefore it can be expected that the effectiveness of both hate and counter speech will be higher once individuals who engage in it organize and act as a group rather than alone.

2 Results

We explore several different macro and micro level measures of effectiveness, providing complementary views from different angles. We make use of a unique situation in Germany where two self-labeling groups engaged in organized online hate and counter speech in 2018. One, called Reconquista Germanica (RG), aimed to spread hate

and misinformation about immigrants, while the other, called Reconquista Internet (RI), tried to actively resist this discourse using organized counter speech. Organizing for RG and related hate groups started in late 2016 [36]. At its peak, RG had between 1,500 and 3,000 active members registered on their Discord server. RI was formed in April 2018, achieving a peak of 62,000 registered users on its Discord server around May 10, 2018. This number quickly decreased to 4-5,000 active members during the first few months, followed by further splintering into smaller independent groups by the end of July 2018. For more information on these two groups see Supplementary Section S1.1.

For the analysis reported here, we used our classification pipeline described in Section 4 and Ref. [29], which not only achieves high accuracy scores (F1 ranging from 0.77 to 0.97) but also aligns well with human judgment to detect hate and counter speech. We used this pipeline to analyze 1,222,240 tweets posted within 181,370 Twitter conversations, or reply trees, that grew in response to tweets of 23 prominent Twitter accounts engaged in political discourse on German Twitter between 2015 and 2018. These accounts comprised prominent news organizations, well-known journalists and bloggers, as well as politicians (see Methods for the full list), all of which were known targets of hate speech. Fig. 1 shows examples of reply trees where each node is a tweet labeled as hate (red), counter (blue), or other (white), with intensity of colors signifying the intensity of the hate score. See Methods for a formal definition of the hate score. Unless stated otherwise, we define hate speech as tweets whose hate score exceeds 0.4 on the scale from -1 (most counter) to 1 (most hate), counter speech as tweets whose hate scores fall below -0.4 on the same scale, and neutral speech as tweets whose score is between -0.4 and 0.4 . Results are robust for other thresholds for hate and counter speech.

2.1 Macro Level Effectiveness Measures

Proportion of hate and counter speech over time

Figure 2 shows a time series of the proportions of counter and hate speech that occurred in our corpus of reply trees (see Section 4.2 for a description of the classification process) over time. Prior to May of 2018, when RI became functional, the relative proportion of both hate and counter speech are largely in a steady state. Roughly 30% of the discourse in our sample of political conversations during this time was hateful and only around 13% was counter speech. However, both of these time series were quite noisy, and there were several very clear deviations reflecting the ebbs and flows of political discourse. Several of these deviations in both hate and counter speech coincided with major events such as large-scale terrorist attacks, political rallies or speeches, and elections. After each of these deviations, however, the proportion of hate and counter speech reverted to the earlier equilibrium, suggesting that these were so-called shock events in an otherwise steady-state system.

After the formation of RI in early spring of 2018, there was a notable change: a decreasing trend in the proportion of hate speech and an increasing trend in the proportion of counter speech. These increasing and decreasing trends continued until approximately September 2018, at which point they more or less stabilized. That point coincided with large alt-right rallies in Chemnitz, Germany, followed by a large counter rally in Berlin. Around that time the proportion of hate and counter speech stabilized at similar levels, with counter speech rising to approximately 21-22% and hate falling to around 25%. Both proportions continued to decline (although very slightly) through the remainder of the studied time period.

While this descriptive analysis of the proportion time series is useful, we further explored these results in two ways: we conducted exploratory change point detection to detect changes in trends over time, as well as a confirmatory analysis of trends before and after the occurrence of RG and RI. The change point detection helped to identify regions of the time series with similar overall trends in proportion of hate and counter speech (see Methods and Refs. [37, 38] for details). These regions are generally separated by “change points” or important events, e.g., terrorists attacks, which shocked the system and impacted the discourse dynamics (sometimes just temporarily). The automatically identified trends are shown in Fig. 2 as straight, thin red and blue lines, accompanied by green vertical lines which signify the change points selected by the algorithm.

The exploratory analysis largely confirmed the descriptive analysis above. For large portions of time, e.g., from the summer of 2016 to a month before the German Federal elections in September of 2017 the trends in hate speech and counter speech were roughly constant. Around the time of the German federal election there is a lot of volatility including a sharply increasing trend in the frequency of hate and a decreasing trend in counter speech. This finding is in line with RG’s primary stated goal of swaying the German elections toward the AfD, a so called “radical right” (“rechtsradikal”) party. After the election and until the formation of RI there is again a constant trend in the proportion of hate and counter speech. Following the formation of RI, there is a noticeable increase in the proportion of counter speech and decrease in the proportion of hate speech, with these trends stabilizing around the time of Chemnitz and Berlin rallies. The analysis also identifies several other periods with decreasing hate and increasing counter speech. However, these periods are effectively relaxation periods following large shock events such as the Islamic terrorist

2.2 Micro Level Effectiveness measures

We begin by analyzing the support that hate and counter speech receives through likes, subsequent replies, and retweets. We then investigate how each hate and counter tweet steers the subsequent discussion in reply trees towards more of the same or towards opposing speech.

Showing support

At the most basic level, participants in Twitter conversations can show their support (or lack thereof) for tweets by affording them their likes and engaging in discussions with them. Figure 3 shows that throughout most of the studied period hate and counter speech received a similar number of likes, but that hate tweets tended to attract longer discussions (measured as the total number of messages in the subtree that a tweet initiates). However, after the emergence of organized counter speech, counter tweets started receiving more likes and attracting longer discussions. The emergence of organized hate speech was not associated with similar changes.




Figure 3: Impact of hate and counter speech messages over time as quantified by the average number of likes and length of conversation they initiate. The emergence of organized counter speech (RI, blue vertical line). Results are for 181,370 reply trees from January 2015 to December 2018. Each data point is a week average and trends are smoothed over a month-long window. The timeline on the x-axis is the same as in other figures but was omitted for space, except for markers of the emergence of RG and RI.

Steering a conversation

A more fine-grained measure of effectiveness is whether individual hate and counter tweets can steer the subsequent conversation. For this analysis, we calculated the change in average discourse from before to after each tweet was posted in a reply tree. We measured the impact I of a tweet as the difference between the average hate score S_h (defined in Eq. 2) of all tweets following and preceding it. More formally, let t_i be the i^{th} tweet that occurred in a tree with N total tweets. Then the impact, or the ability for a tweet t_i to steer a conversations direction, is defined as:

$$I(t_i) = \frac{1}{i-2} \sum_{j=1}^{i-1} S_h(t_j) - \frac{1}{N-(i+1)} \sum_{j=i+1}^N S_h(t_j), \quad (1)$$

where $S_h(t_j)$ is the hate score associated with the j^{th} tweet, t_j . We may refer to t_i as a “focal” tweet when convenient as it is the focus of the computation. We correct the impacts for factors that could have affected it, e.g., discussion length or average hate-score, using a longitudinal mixed linear model (see Methods and Table S1 for details).

Figure 4 shows the results of this analysis. Time is binned into one week segments and shown on the horizontal axis, while the impact of focal tweets are shown on the vertical axis and binned based on their score, going from -1 (most counter) to 1 (most hate). Each square represents the average impact $I(t_i)$ of all tweets occurring in that week with $S_h(t_i) \in [x - 0.05, x + 0.05]$. This plot allows us to study the average ability for tweets with a given hate score to steer conversations in that week. For this analysis, we needed to ensure that a full-blown conversation occurred, so we restricted the sample of reply trees to 82,132 trees that contained at least three tweets (the initial or root tweet, and at least two replies), for a total of 943,822 tweets.

The results shown in Fig. 4 suggest that, in general, hateful tweets were followed by counter speech no matter how intense they are. Counter tweets tend to be followed by hate speech if they are extreme, but by more counter speech when they are moderate. The longitudinal mixed linear model further suggests that tweets near the end of a tree were especially likely to attract more opposition (Table S1). Tweets in larger trees, and in trees that were already biased towards a particular kind of speech, tended to receive more support.

Throughout most of the studied period, political conversations tended to include somewhat more hate than counter speech (Fig. 2), but overall contained a large amount of neutral speech (average hate score across all trees was




Figure 4: Impact of hate and counter focal tweets on subsequent discourse over time. Each colored square represents the average difference between tweets following and preceding a focal tweet posted in a reply tree in a particular week (Eq. 1). Results computed for 82,132 reply trees with at least 3 tweets, over time (horizontal axis), and are binned by hate-score of the focal tweet (horizontal axis; from top to bottom: extreme counter to extreme hate speech). Red (blue) squares signify that the focal tweet was followed by a change in the subsequent conversation towards more hateful (counter) speech, with color saturation corresponding to the size of the difference. Event labels are the same as in Fig. 2.

$S_h = 0.14 \pm 0.048$). However, the large blue areas in Fig. 4 indicate that as conversations progressed, the average scores of the conversations shifted towards more neutral or even the counter side of the speech spectrum. There were only a few exceptional weeks in which discussions tended to be steered toward hate no matter the score of the focal tweet. An example is the week of 2016’s New Year’s Eve when numerous sexual assaults on women in Cologne’s main train station were blamed on Syrian and North African refugees and asylum seekers.

There is a notable difference in discourse dynamics from January 2017 onward, as seen by the transition visible around that time in Fig. 4. This transition corresponds to the time when RG began organizing. Unlike before, hateful conversations in this period did not drift towards counter speech, but stayed neutral or even drifted towards more hateful discourse. This transition is not explained by the classifier picking up on speech used by RG not present before 2017 because, as can be in Fig. 2, the overall proportion of hate speech, as identified by the classifier, is stable around that time. Training data selection bias is also unlikely as shown by Fig. S2. This shift in dynamics is also not explained by any known change in the Twitter interface or algorithms².

Another transition occurred around the time when RI started to organize in the Spring of 2018. Counter speech at that time started to again dampen hate and pull conversations away from hateful rhetoric. This period was followed by several months of backlash during which almost all tweets were associated with more subsequent hate. The backlash likely reflected the broader societal situation in the Summer of 2018 when large alt-right rallies took place throughout Germany, e.g., Chemnitz. Finally, the fall of 2018 was characterized by a return of a more effective counter speech, likely bolstered by large counter rallies occurring in October of 2018, such as “Unteilbar” in Berlin.

Unpacking the pivots

To understand the dynamics unfolding in Fig. 4 in more detail, we unpack how focal tweets pivot the conversations towards either type of speech. We study four possible reactions to each focal tweet: the subsequent discourse can appear to be *supporting* the focal tweet (i.e., counter focal tweet is followed by more counter speech, and vice versa for hate), *opposing* the focal tweet (counter focal tweet is followed by more hate speech), *polarizing* whereby previously neutral discourse becomes more similar to hate or counter speech after the focal tweet, or *ignoring* the focal tweet (discourse remains unchanged before and after the focal tweet). This analysis also reveals the reasons for changes in macro level indicators of hate and counter speech shown in Fig. 2.

Figure 5 shows the relative proportion of different types of reactions that occurred after hate, counter, or neutral focal tweets. In general, hate focal tweets were more often followed by supporting, opposing, or ignoring reactions than counter speech. These trends were relatively stable but could be temporarily altered by notable events, such as the attack on Orlando night club in June 2016, when counter speech received a burst of support. These results suggest that the notable change in discourse dynamics from January 2017 onward, visible in Fig. 4, occurred primarily because neutral speech became more polarized in the direction of hate speech around that time (panel c.). In other words, neutral reply trees started to be steered towards more hateful speech. In addition, hate speech became a bit more supported around that time (panel a.).

This analysis also sheds light on the reasons for relative decrease in frequency of hate speech after the emergence of RI (Fig. 2), as well as for the fact that conversations were more often steered towards counter speech around that time (Fig. 4). Figure 5 suggests that this occurred because of an increase in support for counter tweets (panel a.) and, for a short while, an increased polarization of neutral speech towards counter speech (panel c.). However, as suggested in

²<https://github.com/igorbrigadir/twitter-history/blob/master/changes.csv>

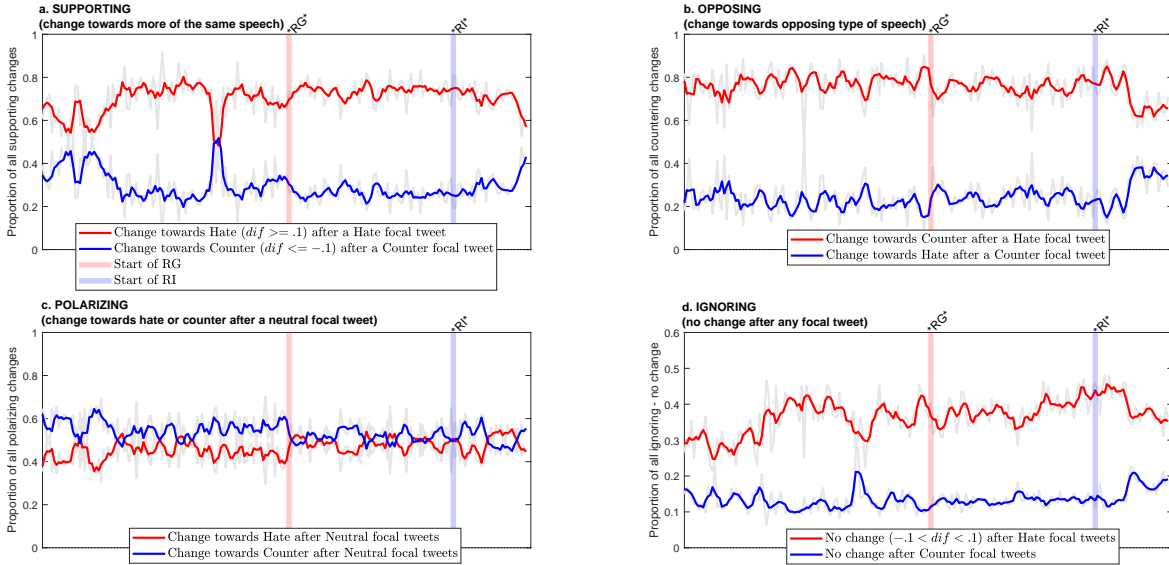


Figure 5: Proportion of focal tweets followed by different reactions, for 82,132 reply trees with at least 3 tweets. Shown are: numbers of focal tweets each week that are followed by a) supporting, or changes towards the same type of speech, b) countering, or changes towards the opposing type of speech, c) polarizing, or changes towards hate or counter after a neutral focal tweet, and d) ignoring the focal tweet. Each data point is a week average and trends are smoothed over a month-long window.

(Fig. 4), a period of hate speech backlash followed the early successes of RI. According to Fig. 5, this occurred mostly because opposing discourse became more frequent after counter speech than after hate speech (panel b.), and neutral speech becomes more polarized towards hate again (panel c.).

3 Discussion

We present the first longitudinal, large-scale study of the dynamics of hate and counter speech in political discussions online. By analyzing hundreds of thousands of Twitter conversations through different lenses at both macro and micro levels, we contribute a nuanced picture of this dynamics. Across a number of different indicators, we find that organized counter speech appears to contribute to a more balanced public discourse. After the emergence of the organized counter group Reconquista Internet (RI) in the late Spring of 2018, the relative frequency of counter speech increased and hate speech decreased (Fig. 2). The number of likes and the length of discussions after counter tweets increased after RI was founded (Fig. 3). Counter speech became more effective in steering conversations when it organized through RI, primarily by providing more support to counter tweets and by steering neutral discourse towards more counter speech as shown in Figs. 4 and 5). While this was met with a backlash initially, the relative frequency of hate speech stabilized into a new, lower proportion of overall speech (Fig. 2). These findings suggest that, like in ‘traditional’ bullying settings, the presence of supporting peers (in this case, other individuals willing to engage in counter speech) motivates people to themselves oppose hate speech and defend its targets.

According to our results, citizens wishing to engage in counter speech would likely increase the effectiveness of their efforts if they organized and participated in discussions in a coordinated way. For example, they could organize around a central platform where members can communicate and strategize (e.g., RI members communicated via Discord server). As an organization they could then steer hateful conversations to a more neutral or even positive ground, by supporting the victim, voicing dissent to hateful positions, or even simply by “liking” counter hateful messages so that they are more broadly visible.

While we see an association of hate and counter speech and subsequent discourse according to a number of indicators, we cannot draw causal conclusions about the effectiveness of organized hate or counter speech alone. German society has been undergoing significant self-examination about its values and political directions throughout the studied period and beyond. This has likely influenced the ongoing discourse in a dynamic interaction with and also beyond organized

counter speech. Regardless, the multifaceted approach we took here suggests that organized counter speech may be a promising strategy in combatting the spread of hate online.

While our team had to go to great lengths to obtain the fully-resolved Twitter conversations used for this analysis, Twitter is now making this kind of analysis easier than ever. With the advent of the second version of the Twitter API researchers can now directly obtain fully resolved conversations that occurred over the past week. This new technology may afford researchers the ability to collect data in additional online contexts, possibly even suitable for analyses of causal influence. We hope this study provides a template for further studies going forward.

4 Methods

4.1 Reply Tree Collection

In total we collected 203,711 conversations (reply trees) that grew in response to tweets of prominent accounts that were either the initiators or the targets of hate speech on German Twitter from 2013 to December of 2018. Originally, we focused our collection on conversations that grew in response to @tagesschau, because it is widely considered the most reliable news source in Germany. This also made it a primary target for hate speech. After an initial period of collection, we chose to add additional Twitter accounts from main-stream news outlets, e.g., @spiegelonline and @faznet. Journalists who were targeted by hate were also added in at a later stage and finally, politicians and other public figures were added as well. In total, these trees contained 1,649,137 tweets, including hate, counter, and neutral speech.

For the analysis reported here we limited this corpus to 181,370 conversations, containing 1,222,240 tweets, which grew in response to tweets from 23 accounts for which we had continuous coverage throughout the period of January 2015 (the beginning of the migrant crisis in Europe) to December 2018. These accounts include prominent news organizations (@diezeit, @derspiegel and @spiegelonline, @faznet, @morgenmagazin, @tagesschau, @tagesthemen, @zdfheute), well-known journalists and bloggers (@annewilltalk, @augstein, @dunjahayali, @janboehm, @jkasek, @maischberger, @nicolediekmann), and politicians (viz., @cem_oezdemir, @c_lindner, @goeringeckardt, @heikomaas, @olafscholz, @regspreecher, @renatekuenast). The number of trees per day from each type of account is shown in Fig. S1.

As the original Twitter API does not provide functionality to collect conversations or “reply trees” in their entirety, we developed custom scrapers (or webcrawlers) to collect these reply trees. There were two kinds of scrapers involved in this process. The first type of scraper systematically identified a random sample of conversations originating from the accounts just discussed. To accomplish this, these scrapers would randomly select one of the accounts above, e.g., @tagesschau, as well as a one month period between January 2013³ and the current date, e.g., 1st of March 2016 to 31st of March 2016. The scraper would then find all top-level urls e.g., https://www.twitter.com/<screen_name>/status/<tweet_id> (where screen name and tweet id refer to the name of the Twitter account and the unique id of the tweet) for tweets by that account in that time period, and stored these in a file. The second class of scrapers would then randomly select a top-level url from this file and manually scrape the resulting conversation in its entirety from the top-level urls raw html. A log of previously scraped reply trees ensured that each tree was only scraped once. The scraper ran continuously from June to December of 2018.

This system was not changed as accounts were added in later phases of the collection process. That is, there was no explicit bias or preference as to which reply trees were scraped in which order. Naturally, accounts that were added to the list later potentially had more “unscraped” trees, and hence, a higher probability of being selected by the reply tree scraper. Conversely, accounts that were added earlier had a higher probability of having a larger number of their trees already scraped. While this could potentially create bias in the resulting sample, we took this into account when selecting both the study period as well as the final list of accounts used in the analysis, ensuring accounts were sampled mostly uniformly across the entire study period. We also investigated the sensitivity of all results to statistical controls for the type of account, and found that all results were similar with or without these controls.

In early 2019, Twitter made fundamental changes to their website’s back-end which made it impossible to collect additional trees in this manner. While our method no longer works for collecting trees, in August 2020 Twitter announced that in the second iteration of the Twitter API, conversations which occurred in the previous week will now be obtainable directly from the API. This means that going forward researchers will be able to do similar analysis to that which we are reporting here without the need for custom software.

³We determined the 1st of January 2013 as our starting point, because AfD was founded shortly after in the same year.

4.2 Classification Pipeline

A detailed description of our classification methods are provided in Ref. [29]. For convenience, we provide a short summary of the most important details of the classification process. To train our classification algorithm, we collected a corpus of more than 9 million relevant tweets originating from known RG accounts (4,689,294 tweets) or RI accounts (4,320,351 tweets). As this data set is only implicitly used in this paper we point the reader to [29] for a more complete description on exactly how these tweets were selected, processed and used. One important detail is that due to limitations of the Twitter API the majority of these training tweets originated during 2018. See Figure S2 for a histogram showing the number of tweets in each category generated per year. The fact that the majority of tweets originated in 2018 is useful, as this was the primary time in which RG and RI interacted and provides the most insight into the language these groups used.

The classification pipeline we used to classify tweets in reply trees consists of two stages ([39]): extraction of features from text and classification using those features. To extract the features, we pre-processed the data and then constructed paragraph embeddings, also known as doc2vec models [40], using the standard gensim implementation [41]. We performed a non-exhaustive parameter sweep following standard practice and the guidelines of [42]. This sweep included the analysis of several doc2vec parameters e.g., maximum distance between current and predicted words, “distributed-memory” vs “distributed bag of words” frameworks, three levels of stop word removal and five different document label types (see [29] for more details). Each version of the doc2vec model was trained on five different but partially overlapping training sets. Each training set included 500,000 randomly selected tweets originating from RG accounts and another 500,000 coming from RI accounts. This produced balanced training sets with 50% hate speech and 50% counter speech. As this data set is only implicitly used in this paper we point the reader to [29], for a more complete description on exactly how these tweets were selected, processed and used.

These trained doc2vec models allowed us to extract features, i.e., infer a feature vector, from a given tweet. To classify each tweet as either hate or counter we then coupled each doc2vec model with a regularized logistic regression classifier, as implemented by scikit-learn [43]. These logistic regression functions were trained on the same training set as the corresponding doc2vec model.

By pairing a trained doc2vec model and a logistic regression function, we were able to assign a probability $p_h(t)$ to each tweet that it belongs to either the hate or counter class. We recoded this probability to a hate score $S_h(t)$ ranging from -1 to 1, using

$$S_h(t) = 2 * p_h(t) - 1 \quad (2)$$

where negative values mean that a tweet corresponds to counter speech and positive values mean that it corresponds to hate speech. For the results, reported here we used an ensemble learning approach where 25 independent doc2vec, logistic regression pairings voted on each tweet in a given tree. The tweets final hate score $S_h(t)$ was then defined as the average hate score assigned to that tweet across all 25 pairings.

For final classification of tweets, we used a confidence voting system with thresholding to assign each tweet a label. For the majority of the analyses reported in this paper (unless stated otherwise), if $S_h(t) \geq 0.4$ was labeled hate, and if $S_h(t) \leq -0.4$ then t was labeled counter. If $-0.4 \leq S_h(t) \leq 0.4$ then the tweet was labeled neutral. These scores effectively mean that the ensemble of classifiers was at least 70% confident in labeling a tweet as either hate or counter.

Depending on the threshold being used, this system achieved F1 scores ranging from 0.755 to 0.97 on balanced test sets containing equal proportion of hate and counter speech. For the threshold used in this paper the classification pipeline achieved F1 scores of 0.875. This accuracy exceeds previous studies that used smaller unbalanced data sets and achieved F1 scores ranging from 0.49 to 0.77 [25, 26, 44]. See [29] for further details about this classification system.

4.3 Statistical Analysis

To confirm whether counter speech had an effect on discourse dynamics, one would ideally like to perform some flavor of statistical causal inference analysis. However, due to a variety of conflating factors, the reply tree time series presented in this paper are nearly impossible to analyze from a causal perspective. Specifically, because of unquantifiable outside influences it would not be appropriate to causal inference methods such as Granger Causality [45] or Transfer Entropy [46], and a lack of evidence for determinism and stationarity in our trends make state-space [47] based causal inference, e.g., Convergent Cross Mapping [48], not applicable.

While we could not perform rigorous causal inference due to fundamental limitations of the study system, we nevertheless put effort into conducting a principled study of the changes in these dynamics over time—even if such an analysis could only bring out correlation structure and suggest changes in descriptive statistics such as overall trends.

In particular, to investigate changes in the proportions of hate and counter speech over time and in association with different events (Fig. 2), we conducted both exploratory and confirmatory statistical analyses. Exploratory change

point analyses were conducted using the “findchangepts” algorithm available in Matlab (Refs. [37, 38]). The algorithm searched for changes in linear trends, always using a standard deviation of the empirical trend as the threshold value. There are many change point algorithms we could have used for this analysis (see [38] for a review) which identify changes in some summary statistic, e.g., mean, variance or information production. However, we were interested in studying changes in overall *trends* in discourse which made [37] an ideal approach. For a confirmatory analysis we tested differences in trends before and after occurrence of Reconquista Germanica and Reconquista Internet.

To investigate the effectiveness of different types of speech to steer a conversation depending on its hate score (Fig. 4), we corrected the impact of focal tweets using a longitudinal mixed linear model. In this model, focal tweet impact was predicted by its score and a host of other factors that could have affected the impact in addition to the score. The focal tweet’s position in the tree, tree size, average hate score of the tree, week, and type of Twitter account that started the reply tree (news outlets, politicians, or journalists/bloggers) were included as fixed effects, and variations by tree and week, tree size were allowed by including them as random effects (see Table S1. This correction did not change original results much (see top panel of Fig. S3 for uncorrected results). Note that each square in (Fig. 4) represents the average of all focal tweets with the particular hate score posted in the specific week. Different squares include different number of such tweets, as shown in bottom panel of Fig. S3, and different number of trees. These and other factors has been statistically accounted for in using the model described above.

Acknowledgment

J.G. was partially supported by an Omidyar and an Applied Complexity Fellowship at the Santa Fe Institute. J.-G.Y. was supported by a James S. McDonnell Foundation Postdoctoral Fellowship Award. L.H.-D. and J.-G.Y. were supported by Google Open Source under the Open-Source Complex Ecosystems And Networks (OCEAN) project. M.G. was partially supported by NSF-DRMS 1757211. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Google Open Source or the NSF.

References

- [1] Chara Bakalis. *Cyberhate: An issue of continued concern for the Council of Europe’s Anti-racism Commission*. Council of Europe, 2015.
- [2] Catherine Blaya. Cyberhate: A review and content analysis of intervention strategies. *Aggression and violent behavior*, 45:163–172, 2019.
- [3] Iginio Gagliardone, Danit Gal, Thiago Alves, and Gabriela Martinez. *Countering online hate speech*. Unesco Publishing, 2015.
- [4] James Hawdon, Atte Oksanen, and Pekka Räsänen. Exposure to online hate in four nations: A cross-national consideration. *Deviant Behavior*, 38(3):254–266, 2017.
- [5] Karsten Müller and Carlo Schwarz. Fanning the flames of hate: Social media and hate crime. *Available at SSRN 3082972*, 2019.
- [6] Atte Oksanen, Markus Kaakinen, Jaana Minkkinen, Pekka Räsänen, Bernard Enjolras, and Kari Steen-Johnsen. Perceived societal fear and cyberhate after the november 2015 paris terrorist attacks. *Terrorism and Political Violence*, pages 1–20, 2018.
- [7] Anne Weber. *Manual on hate speech*. Council Of Europe, 2009.
- [8] Youtube: Hate speech policy. <https://support.google.com/youtube/answer/2801939>. Accessed: 2020-02-28.
- [9] Twitter: Hateful conduct policy. <https://help.twitter.com/en/rules-and-policies/hateful-conduct-policy>. Accessed: 2020-02-28.
- [10] Facebook: Hate speech. https://www.facebook.com/communitystandards/hate_speech. Accessed: 2020-02-28.
- [11] Eshwar Chandrasekharan, Umashanthi Pavalanathan, Anirudh Srinivasan, Adam Glynn, Jacob Eisenstein, and Eric Gilbert. You can’t stay here: The efficacy of reddit’s 2015 ban examined through hate speech. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW):1–22, 2017.
- [12] S Benesch, D Ruths, KP Dillon, HM Saleem, and L Wright. Considerations for successful counterspeech, 2016.
- [13] Diana Rieger, Josephine B Schmitt, and Lena Frischlich. Hate and counter-voices in the internet: Introduction to the special issue. *SCM Studies in Communication and Media*, 7(4):459–472, 2018.
- [14] Sebastian Wachs, Michelle F Wright, Ruthaychonnee Sittichai, Ritu Singh, Ramakrishna Biswal, Eun-mee Kim, Soeun Yang, Manuel Gámez-Guadix, Carmen Almendros, Katerina Flora, et al. Associations between witnessing and perpetrating online hate in eight countries: the buffering effects of problem-focused coping. *International journal of environmental research and public health*, 16(20):3992, 2019.
- [15] Marc Ziegele, Pablo Jost, Marike Bormann, and Dominique Heinbach. Journalistic counter-voices in comment sections: Patterns, determinants, and potential consequences of interactive moderation of uncivil user comments. *SCM Studies in Communication and Media*, 7(4):525–554, 2018.

- [16] Jürgen Habermas. *Between facts and norms: Contributions to a discourse theory of law and democracy*. John Wiley & Sons, 2015.
- [17] Éloi Brassard-Gourdeau and Richard Khoury. Impact of sentiment detection to recognize toxic and subversive online comments. *arXiv preprint arXiv:1812.01704*, 2018.
- [18] Pete Burnap, Omer F Rana, Nick Avis, Matthew Williams, William Housley, Adam Edwards, Jeffrey Morgan, and Luke Sloan. Detecting tension in online communities with computational Twitter analysis. *Technological Forecasting and Social Change*, 95:96–108, 2015.
- [19] Pete Burnap and Matthew L Williams. Us and them: identifying cyber hate on Twitter across multiple protected characteristics. *EPJ Data science*, 5(1):11, 2016.
- [20] Sean MacAvaney, Hao-Ren Yao, Eugene Yang, Katina Russell, Nazli Goharian, and Ophir Frieder. Hate speech detection: Challenges and solutions. *PLoS one*, 14(8), 2019.
- [21] Manoel Horta Ribeiro, Pedro H Calais, Yuri A Santos, Virgílio AF Almeida, and Wagner Meira Jr. Characterizing and detecting hateful users on Twitter. In *Twelfth international AAAI conference on web and social media*, 2018.
- [22] Ziqi Zhang and Lei Luo. Hate speech detection: A solved problem? the challenging case of long tail on Twitter. *Semantic Web*, 10(5):925–945, 2019.
- [23] Steven Zimmerman, Udo Kruschwitz, and Chris Fox. Improving hate speech detection with deep learning ensembles. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, 2018.
- [24] Hannah Gaffney, David P Farrington, Dorothy L Espelage, and Maria M Ttofi. Are cyberbullying intervention and prevention programs effective? a systematic and meta-analytical review. *Aggression and violent behavior*, 45:134–153, 2019.
- [25] Binny Mathew, Navish Kumar, Pawan Goyal, and Animesh Mukherjee. Analyzing the hate and counter speech accounts on Twitter. *arXiv preprint arXiv:1812.02712*, 2018.
- [26] Binny Mathew, Punyajoy Saha, Hardik Tharad, Subham Rajgaria, Prajwal Singhanian, Suman Kalyan Maity, Pawan Goyal, and Animesh Mukherjee. Thou shalt not hate: Countering online hate speech. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 13, pages 369–380, 2019.
- [27] Natalie Jomini Stroud, Joshua M Scacco, Ashley Muddiman, and Alexander L Curry. Changing deliberative norms on news organizations’ facebook sites. *Journal of Computer-Mediated Communication*, 20(2):188–203, 2015.
- [28] Lucas Wright, Derek Ruths, Kelly P Dillon, Haji Mohammad Saleem, and Susan Benesch. Vectors for counterspeech on Twitter. In *Proceedings of the first workshop on abusive language online*, pages 57–62, 2017.
- [29] Joshua Garland, Keyan Ghazi-Zahedi, Jean-Gabriel Young, Laurent Hébert-Dufresne, and Mirta Galesic. Countering hate on social media: Large scale classification of hate and counter speech. *arXiv:2006.01974*, 2020.
- [30] Marc Ziegele and Pablo B Jost. Not funny? the effects of factual versus sarcastic journalistic responses to uncivil user comments. *Communication research*, page 0093650216671854, 2016.
- [31] Christina Salmivalli. Participant roles in bullying: How can peer bystanders be utilized in interventions? *Theory Into Practice*, 53(4):286–292, 2014.
- [32] Gianluca Gini, Paolo Albiero, Beatrice Benelli, and Gianmarco Altoe. Determinants of adolescents’ active defending and passive bystanding behavior in bullying. *Journal of adolescence*, 31(1):93–105, 2008.
- [33] Christina Salmivalli, Marinus Voeten, and Elisa Poskiparta. Bystanders matter: Associations between reinforcing, defending, and the frequency of bullying behavior in classrooms. *Journal of Clinical Child & Adolescent Psychology*, 40(5):668–676, 2011.
- [34] Amalia Álvarez-Benjumea and Fabian Winter. Normative change and culture of hate: An experiment in online environments. *European Sociological Review*, 34(3):223–237, 2018.
- [35] J Nathan Matias. Preventing harassment and increasing group participation through social norms in 2,190 online science discussions. *Proceedings of the National Academy of Sciences*, 116(20):9785–9789, 2019.
- [36] Jacob Davey and Julia Ebner. The fringe insurgency. connectivity, convergence and mainstreaming of the extreme right. *Institute for Strategic Dialogue* (<http://www.isdglobal.org/wp-content/uploads/2017/10/The-Fringe-Insurgency-221017.pdf>), 2017.
- [37] Rebecca Killick, Paul Fearnhead, and Idris A Eckley. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, 107(500):1590–1598, 2012.
- [38] Samaneh Aminikhanghahi and Diane J Cook. A survey of methods for time series change point detection. *Knowledge and information systems*, 51(2):339–367, 2017.
- [39] Anna Schmidt and Michael Wiegand. A survey on hate speech detection using natural language processing. In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, pages 1–10, 2017.
- [40] Quoc Le and Tomas Mikolov. Distributed representations of sentences and documents. In *International conference on machine learning*, pages 1188–1196, 2014.

- [41] Radim Řehůřek and Petr Sojka. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta, May 2010. ELRA. <http://is.muni.cz/publication/884893/en>.
- [42] Jey Han Lau and Timothy Baldwin. An empirical evaluation of doc2vec with practical insights into document embedding generation. *arXiv preprint arXiv:1607.05368*, 2016.
- [43] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [44] Caleb Ziems, Bing He, Sandeep Soni, and Srijan Kumar. Racism is a virus: Anti-asian hate and counterhate in social media during the covid-19 crisis. *arXiv preprint arXiv:2005.12423*, 2020.
- [45] Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pages 424–438, 1969.
- [46] Thomas Schreiber. Measuring information transfer. *Physical review letters*, 85(2):461, 2000.
- [47] Floris Takens. Detecting strange attractors in turbulence. In *Dynamical systems and turbulence, Warwick 1980*, pages 366–381. Springer, 1981.
- [48] George Sugihara, Robert May, Hao Ye, Chih-hao Hsieh, Ethan Deyle, Michael Fogarty, and Stephan Munch. Detecting causality in complex ecosystems. *science*, 338(6106):496–500, 2012.

S1 Supplementary Information

S1.1 Background on Reconquista Germanica and Reconquista Internet (RI)

Reconquista Germanica (RG) officially formed two months prior to the German federal election with the purpose to influence the election to benefit Germany’s far-right party The Alternative für Deutschland (AfD). However, some sources claim that RG began organizing in late 2016 and were active as early as January 2017 [36]. RG was organized hierarchically, and prospects had to prove their political views in a series of online interviews, before they were granted access to the Discord server. Once granted, people could climb the ranks by proving themselves in coordinated actions, which were given in daily orders to the entire group. Actions were then coordinated between multiple social media outlets e.g, Twitter, Facebook, Twitch, YouTube, etc. Each participants was usually responsible for multiple accounts. It is unclear how many active members RG had and how many accounts they controlled. Estimates are about 1500 active members controlling 10.000 - 20.000 accounts on various social media platforms.

According to publicly available leaks⁴ of Reconquista Internet (RI) during their initial time, RI was organized in a loose hierarchical structure. A core group around the founder Jan Böhmermann proposed challenges or so-called *fluffy attacks*. An example of such a fluffy attack was the challenge to engage AfD politicians in a discussion and make them state a variation of the first article of German’s federal law. The variation was something along the following statement “The dignity of every human shall be inviolable.” instead of “Human dignity shall be inviolable.”

Similar to RG, RI was organized into groups which formed around specific social media platform, i.e., there was a group for Twitch, Facebook, Twitter, etc. Each group had two speakers, who organized activities within their group. To communicate with one another, call for help etc, several Discord channels were maintained, e.g., a channel for sharing links to tweets which were selected for engagement, under attack etc.

Both RG and RI were active on many platforms e.g., Twitter, YouTube and Facebook. Here we focus on their Twitter presence, which enables us to study the structure of the resulting conversations.

S1.2 Supplementary Figures




Figure S1: Data set: Trees per day, by type of account, for a total of 181,370 trees used in analyses.

S1.3 Supplementary Tables

⁴<https://requileaks.github.io>




Figure S2: Data set: Number of tweets per year that were collected to train the classification pipeline. Sample counts of training examples for each class collected per year.




Figure S3: Top panel: Uncorrected results for Fig. 4. Bottom panel: Number of focal tweets in each category of focal tweet score.

	Estimate	Lower Bound	Upper Bound
Fixed effects			
Intercept	-0.0994	-0.1179	-0.0810
$S_h(t_i)$	-0.0182	-0.0265	-0.0099
Focal tweet position $\text{pos}(t_i)$	0.0001	0.0001	0.0002
Tree size	0.0002	0.0001	0.0003
Average tree hate score	-0.0097	-0.0311	0.0116
Week w	0.0002	0.0002	0.0003
Politicians vs. news outlets	0.0690	0.0597	0.0784
Journalists/bloggers vs. news outlets	0.0846	0.0758	0.0934
$S_h(t_i)$ x $\text{pos}(t_i)$	-0.0005	-0.0005	-0.0005
$S_h(t_i)$ x Tree Size	0.0002	0.0002	0.0003
$S_h(t_i)$ x Average tree hate score	0.0165	0.0058	0.0273
$S_h(t_i)$ x w	0.0000	0.0000	0.0001
$S_h(t_i)$ x Politicians	0.0010	-0.0021	0.0041
$S_h(t_i)$ x Journalists	0.0009	-0.0023	0.0041
Random effects			
Tree	0.3515	0.3495	0.3536
Week	0.0152	0.0117	0.0196

Table S1: Longitudinal mixed linear model of differences due to focal tweets, implemented by matlab fitlme procedure. More complicated models did not produce better fit indices. The results suggest that authors of both hate and counter speech tend to show clustering behavior. Focal tweets in trees that are overall biased towards the same type of speech, larger trees, and trees started by politicians and journalists as opposed to official news outlets, tend to attract more of the same speech. In other words, in such trees hate was somewhat more likely to attract hate, and counter to attract counter speech. Accounts of politicians and journalists were overall more likely to receive hateful responses compared to traditional news outlets. Tweets near the end of the tree tend to be followed by more opposition (which in turn possibly contributes to ending of the discussion).