Online hate speech and emotions on Twitter.

A case study of Greta Thunberg at the UN Climate Change Conference COP25 in 2019

Arce-García, Sergio. Universidad Internacional de La Rioja - Logroño (La Rioja - Spain)
sergio.arce@unir.net (Corresponding author)
https://orcid.org/0000-0003-0578-9787

Díaz-Campo, Jesús. Universidad Internacional de La Rioja - Logroño (La Rioja - Spain)
jesus.diaz@unir.net
https://orcid.org/0000-0001-5014-8749

Cambronero-Saiz, Belén. Universidad Internacional de La Rioja - Logroño (La Rioja - Spain)
belen.cambronero@unir.net
https://orcid.org/0000-0001-5273-9534

ABSTRACT

The presence of environmental activist Greta Thunberg at the UN Climate Change Conference COP25 in 2019 prompted reactions on social media, which grew exponentially after she was named Time Magazine's Person of the Year 2019 and even more so after then-president of the United States Donald Trump tweeted his reaction to her accolade. An analysis of 1,395,054 tweets gathered between November and December 2019 through R, network theory techniques, machine learning and natural language processing showed how messages sparking hatred and intense emotions generate posts, mainly negative ones, that subsequently serve as catalysts. The results also demonstrate the relevance of the bubble filter and echo chamber theories and the fact that hate springs from a range of sentiments depending on each participant group.

KEY WORDS

Greta Thunberg, Climate Summit, social media, hate, machine learning, sentiment.

STATEMENT AND DECLARATIONS

We wish to confirm that there are no known conflicts of interest associated with this publication and there has been no significant financial support for this work that could have influenced its outcome. We confirm that the manuscript has been read and approved by all named authors and that there are no other persons who satisfied the criteria for authorship but are not listed.

We further confirm that the order of authors listed in the manuscript has been approved by all of us.

THEORETICAL FRAMEWORK

Social media communication has become increasingly relevant in shaping public opinion (Bode 2016; Gil de Zúñiga et al. 2017). Conversations and debates traditionally hosted on news outlets now take place or are amplified on social media networks (Jansson and Lindell 2015; García-Perdomo et al. 2018). This phenomenon, in which permanent connectivity, multiscreen or mobile devices play a fundamental role (Picone et al. 2015), is particularly pertinent to Twitter, a social media platform that lends itself to this form of debate due to, among other things, its short-messaging style (Wasike 2017; Díaz-Campo et al. 2021).

Thus, social media, and particularly Twitter, are currently a scenario for the rapid dissemination of all manner of messages (Oltmann et al. 2020; Segado-Boj et al. 2020). Speed of dissemination prevails over rigour or verification (Urman 2019), generating a scenario characterised by polarization and fragmentation (Sunstein 2017) where highly active users seek and find messages that strengthen their beliefs. With pluralism and the debate of ideas in decline (Serrano-Puche et al. 2020), each social group seeks to adapt messages to their own interests (Evolvi 2017).

The dynamic of social media networks, which allows users to become active broadcasters, causes these groups to expand (Karlsen 2015) and transforms the concept of opinion leaders, or influencers, who connect with and persuade a wide range of audiences. This generates various processes such as filter bubbles (Pariser 2011) through which users only receive information according to their interests and preferences. However, such information only reinforces their own beliefs and opinions, thus cutting them off from opposing ideas. Similarly, the echo chamber (Colleoni et al. 2014), where opinion groups share information and establish social cohesion (Duffy and Ling 2020), makes these particularly active users feel they are helping to spread valuable and interesting content to society. However, they are communicating with a limited audience, their own circle, and like-minded users (Agur and Gan 2021).

One of the most widespread techniques used to measure all these aspects is sentiment and emotion analysis, which, via machine-learning techniques and algorithms, retrieves texts and classifies them around a series of positive and negative emotions (Tausczik and Pennebaker 2010). These emotions are not necessarily contained in the message itself, but they can be triggered by it (Chang 2019) and be used to anticipate phenomena such as crises (Arcila-Calderón et al. 2017).

Among the most prevalent emotions found on Twitter is hate (Paz et al. 2020). This sentiment, which derives from the basic emotions (Plutchik 1980), is so prevalent on social media that it has led to the coining of expressions like "cyberhate" or "hate speech" to refer to messages, generally propagandistic and with obvious violent connotations, targeting specific groups of people due to a perceived identity or difference (race, gender, ideology, etc.) (Garland and Chakraborti 2012; Williams 2021).

Previous studies have shown that when messages are read, they can trigger emotions, even if they contain no emotional content per se (Duncombe 2019), or that the emotions generated upon reading the message are more determinant than the content itself (Berger and Milkman 2012; Weeks and Holbert 2013; Keib et al. 2018; Arce-García, et al. 2020). Other studies suggest a greater tendency to share news that generates negative emotions such as fear or anger (Segado-Boj et al. 2020).

UN Climate Change Conference COP25

The UN Climate Change Conference COP25, held between 2 December and 15 December 2019 in Madrid, was attended by more than 25,000 representatives from some 250 countries. The main aim of the meeting was to seek accords and commitments to tackle the effects of climate change, with a particular focus on cutting greenhouse gas emissions. Indeed, the draft text finally approved stressed the "urgent need" for countries to reduce these emissions.

With some 1,500 journalists accredited to cover the COP 25, the event generated considerable interest on social media. One of the most visible figures at the event was the Swedish environmental activist Greta Thunberg, a prime example of an opinion leader or influencer. Besides attending the "Social Summit for Climate in Madrid (Beyond COP25: People for Climate), a "parallel event counter-event" that brought together a host of social and environmental organizations, Thunberg also took part in a "Climate March" on 6 December.

Given Thunberg's controversial public reputation, we believe that a study on social media activity around the event is pertinent for analyzing the presence and incidence of hate speech. This research provides a way of dealing with hate and its association with emotions caused by an event, depending on each group detected, with the help of machine learning techniques applied to the field of communication and social networks.

This research aims to identify the movements on the social network Twitter around the visit of Greta Thunberg during the COP 25 in Madrid in 2019. It argues that the visit of this activist shifted the public debate to aspects not directly related to the conference, but to polarised spaces between well-differentiated groups.

Methodology

The research frames the considerable impact the environmental activist Greta Thunberg had on the UN Climate Change Conference COP25 in Madrid from 2 December to 13 under the presidency of the Government of Chile and with logistical support of the Spanish Government. In order to grasp Thunberg's influence on both the event and world public opinion, this study tracked all tweets mentioning @GretaThunberg sent during the period leading up to and shortly after the Conference. The inclusion criteria were therefore to collect all messages, tweets and retweets in any language that mentioned the climate activist in their message, first eliminating the possibility of duplication during the analysis period.

The messages were gathered between 11/25/2019 at 09:19 GMT and 12/17/2019 at 11:15 GMT through the statistical software R via RStudio version 1.2.5033 via the RTweet library (Kearney 2019) with the help of the Twitter application-programing interface (API) for developers. An overload in messages posted worldwide on 11 December led to input problems, which exceeded the number of tweets per period of the Twitter API 1.1, meant that no data were obtained for some 30 hours. This issue was resolved by extracting the unsuccessful capture of the time area to Twitter API 2.0 for business and academic work via the R library AcademicTwitteR and accredited academic account (Barrie and Chun-ting 2021). A total of 1,395,054 tweets were captured and then processed as follows:

- Graphic analysis through network theory in communication (Barabasi 2016), by the program Gephi version 0.9.5, establishes connections between user nodes. The retweet representation algorithm ForceAtlas2 was used (Jacomy et al. 2014) to establish a cluster analysis to determine groups of influence (Chen et al. 2020) through the optimization methodology following Louvain's algorithm (Blondel et al. 2008).

- Statistical analysis using R and determination of bot behaviour probability through Kearney's Tweetbotornot algorithm (2018). This algorithm, which uses machine learning techniques, is based on the history of the account's last 100 messages, its friends and followers, profile and photo, among other available data, and is used with a 93.8% probability of success.

- Natural Language Processing (NLP) analysis and quantification of polarity and emotions based on the National Research Council of Canada (NRC) lexicon, with more than 14,000 words in its 0.92 version for each language. In this study, English and Spanish (Mohammad and Turney 2010, 2013) were used along with the Syuzhet R-library (Jockers 2017). Emotions were labelled according to the primary classification of emotions by Plutchik (1980) and subsequently by Ekman (2003), based on the universality of emotions and facial expressions, these being: anticipation (expression of rational thought), disgust, fear, joy, sadness, surprise, and confidence, which were applied to the discourse in the tweets (Sauter et al. 2010). Each text had an associated valence (positive or negative) and/or intensity that was enhanced or diminished by the words around it (Swati et al. 2015), ultimately calculating the total value of the whole message. In this way, the polarity values offer values of different intensity in the positive or negative sense, and are neutral when their value is equal to zero.

- Analysis and quantification of hate using PLN, supported by the 2018 version of Hurtlex lexicon in English and Spanish, with 8,228 and 5,007 associated words respectively (Bassignana et al. 2018). The R Syuzhet library was used to detect and quantify the lexical corpus. As English was the most common language in all the groups identified, except for two in Spanish, each cluster was studied in its main language only.

- Regression trees analysis, which can predict the behaviour of a variable through machine-learning (Lantz 2019). Although different emotions can determine hate, the term was treated as a dependent variable on the basic emotions and polarity, which were independent, and other variables of position in the network such as eigenvector, betweenness centrality or closeness centrality. For this purpose, the classification and regression tree (CART) algorithm (Breiman et al. 1984) was used through the "rpart" package in R (https://cran.r-project.org/web/packages/rpart/index.html). The prediction is tested by estimating the correlation values between the actual and predicted values and the mean absolute error (MAE).

- Closeness centrality analysis in relation to hate generated through R. This closeness value is a metric of the nodes representing their mutual relationship, where those farthest from the centre show values close to 0 and are closely linked to others at 1 (Hansen et al. 2011). This provides information on the time it takes messages to be disseminated to the network according to their closeness to other nodes (Rodrigues 2021).

A PC with an i7 processor and 32 GB of RAM running Linux was used to obtain and process the calculations for all these techniques; there was no need to compute in the cloud.

Results

The data obtained that referenced @GretaThunberg before, during, and after the UN Climate Change Conference totalled 1,395,054 tweets, of which 412,898 were direct messages and 982,156 retweets (RT). Of these, 1,112,887 were in English (79.77%), 74,688 in Spanish (5.35%), 57,652 undetermined (4.13%), 36,691 in Portuguese (2.63%), 25,591 in German (1.83%), 18,203 in Haitian (although they were really in English, 1.30%), 16,379 in French (1.17%), 6,331 in Swedish (0.45%), and 3,103 in Dutch (0.22%).

Figure 1 shows the number of tweets collected over the study period, showing an increase in tweets in early December, followed by a surge of tweets around @GretaThunberg on December 11 and 12, when two relevant events occurred:

- Time Magazine named Greta Thunberg its "Person of the Year" and featured her on its cover on 11 December (Alter et al. 2019).
- US President Donald Trump taunted Thunberg's accolade of Time Magazine's Person of the Year, with the tweet, "So ridiculous. Greta must work on her Anger Management problem, then go to a good old fashioned movie with a friend! Chill Greta, Chill! https://t.co/M8ZtS8okzE - Donald J. Trump (@realDonaldTrump)" (Noor 2019) Almost almost a year later, on 5 November 2020, when Trump lost the US presidential election, Greta Thunberg tweeted the same message to Trump, (Reuters 2020).

Figure 1: Number of Tweets around @GretaThunberg. Source: Own elaboration

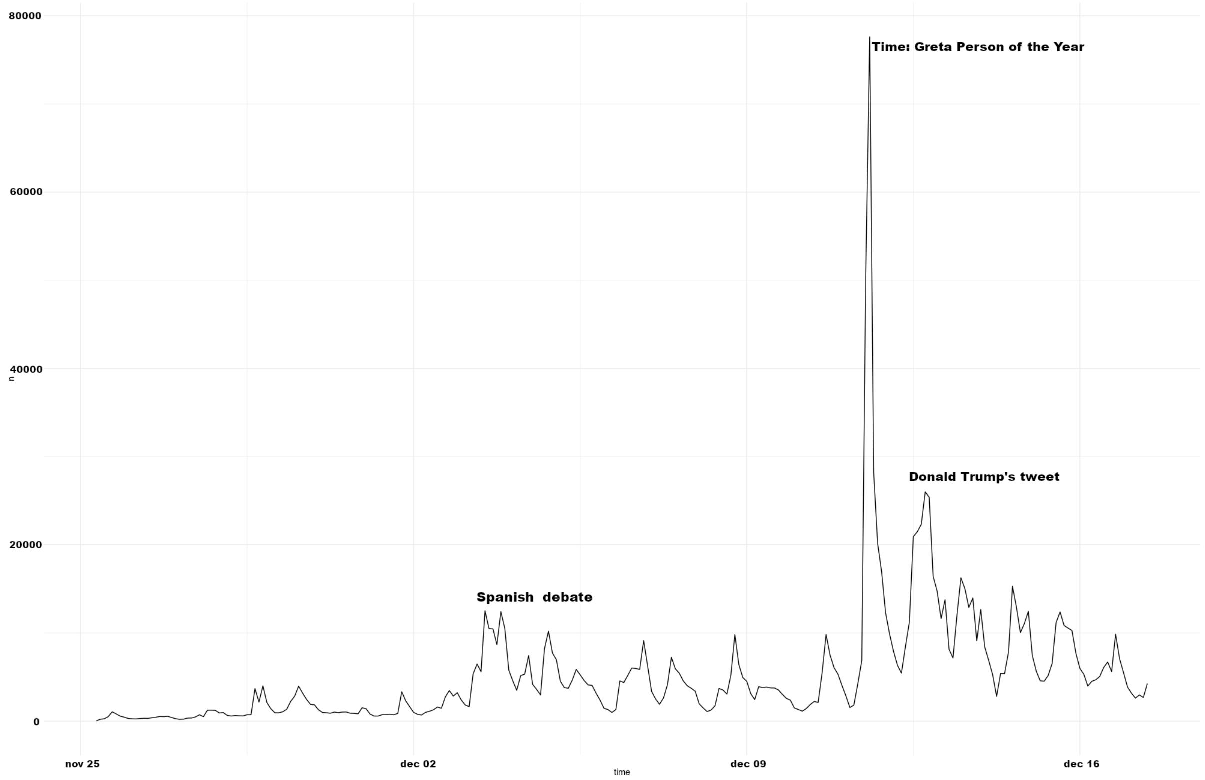

A network-theory analysis of retweets (RT) established the relationships and the formation of opinion groups or clusters, shown in Figure 2. The average degree or mean number of tweets sent per user was calculated at 1.443. The modularity of the network was 0.564, a good value for social media, with the formation of 6,347 opinion groups, but where the top ten accounted for 86.04% of the total retweet traffic. The following groups in order of the number of tweets were (the number of the group comes from the cluster identified in the operation of the algorithm):

- Group 0: Greta and followers (38.06% of traffic), shown in pink in Figure 2. The primary account followed was @GretaThunberg. In English.
- Group 723: US Democratic politicians and media (10.9%), in pale green. Top influencer accounts were @HillaryClinton, @JohnKerry, @JoeBiden, @RepMaxineWatersen or CNN journalist @ananavarro. In English.
- Group 1,173: Time Magazine (10.26%), in blue, around the magazine's official account. In English.

- Group 389: Anti-Thunberg group and media (7.02%), in brown. The main hubs were: Fox News journalists (@SaraCarterDC, top influencer of the group) and Republican influencers such as @RaheemKassan or @MaximeBernier, or accounts, such as @1776AmericaUSA, later suspended for breaking Twitter rules. In English.

- Group 24: Followers of Michelle Obama (5.86%), in orange. In English.

- Group 18: Followers of @MikeHudema environmental activist (3.63%), in pink, supporting Greta. In English.

- Group 1,224: Followers of actor @LeoDicaprio, with a considerable influence in Brazilian networks and its media (3.2%), in dark green. Supporting Greta. In English and Portuguese.

- Group 1,150: Followers of @Joshuawongcf, Hong Kong activist (3.13%), in light pink. In English.

- Group 285: Spanish left-wing group and related media (2.52%), in yellow. In Spanish.

- Group 225: Spanish far-right group (1.46%), in red, in Spanish.

Centrality (eigenvector equal to 1) was determined around @GretaThunberg's account, with all groups supporting her discourse around it and many connections between groups. The groups supporting Thunberg in languages other than English, mainly Portuguese and Spanish, are also close to the main core but have less interaction with other groups due to language barriers. Among the Democratic representatives from the United States, the group following Michelle Obama, the former First Lady of the United States, stands alone, despite being highly interconnected.

Thunberg's critics are only linked via the Hong Kong and Time Magazine groups.

All groups have major influencers who tweet and are followed and retweeted by their users, but the anti-Thunberg groups, especially the Spanish group, have many unknown nano-influencers.

Figure 2: Network graph around @GretaThunberg. Source: Own elaboration

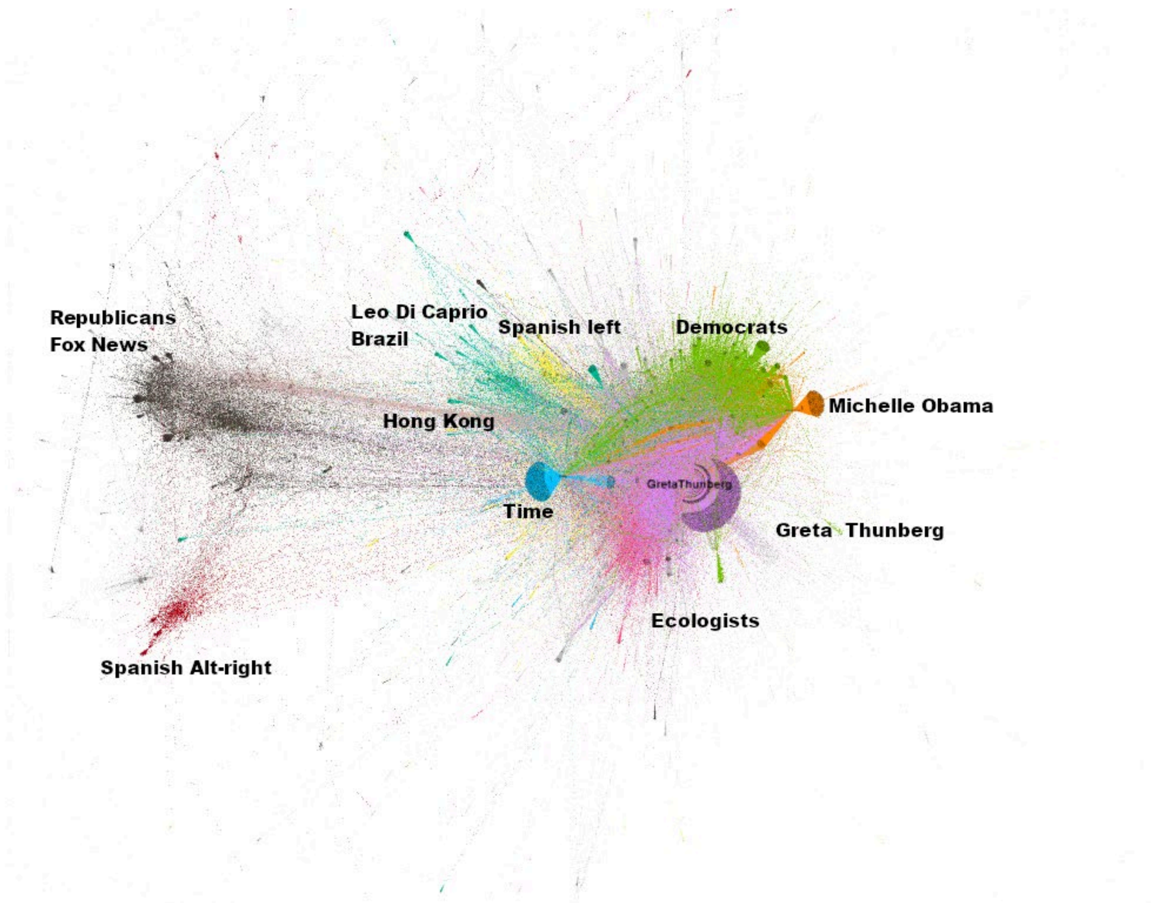

There is no uniform distribution of messages in each group during the analyzed period. The boxplot diagrams in Figure 3 show that the politically antagonistic Spanish-speaking groups were first, mainly the one linked to the Spanish far right (225), immediately followed by left-wing political groups from the same country. The anti-Thunberg group (389), which was the first to reach 25% of tweets in English, remained the most active for the longest, with half of its messages tweeted on the days of Time Magazine's Person of the Year award and Donald Trump's tweet. The groups following Greta Thunberg and other environmental activists engaged more heavily in the run-up to the Conference, while the rest of the groups focused on the days of the Time Magazine cover and the US president's tweet in response to it.

Figure 3: Box plot showing the tweets posted over the period under study. Source: Own elaboration

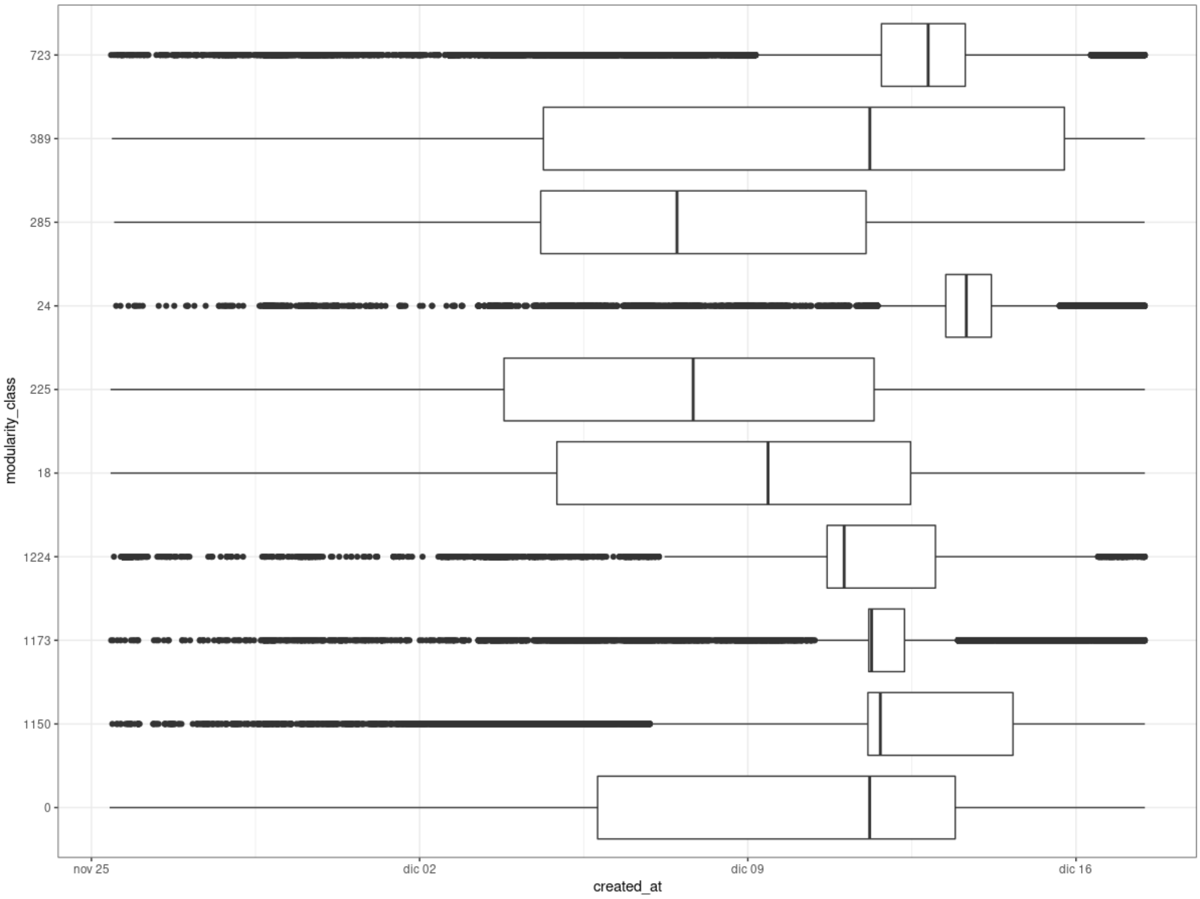

Emotion and valence study

Each group's positive and negative polarity and intensity analysis (Figure 4) shows how almost all the groups have a mean, indicated with a blue line, close to neutrality when considering the totality of messages. However, Michelle Obama's groups (24), the Hong Kong activists (1150) and Greta's own group (0) increased in mean positivity. Of particular note is the large amount of highly polarized content (above 3 or below -3) in quite a few groups, which increases after Thunberg's appearance on the cover of Time Magazine and after Donald Trump's tweet. Towards the end of the period under study, the emergence of many more tweets with a particularly negative polarity in groups 0, 18, 723 and 285 makes this evident. These groups coincide with the followers of Greta Thunberg and environmental activist @MikeHudema on the one hand and those attacking Greta Thunberg in English and Spanish on the other.

Figure 4: Polarity analysis from 11/25/2019 to 12/17/2019. Source: Own elaboration

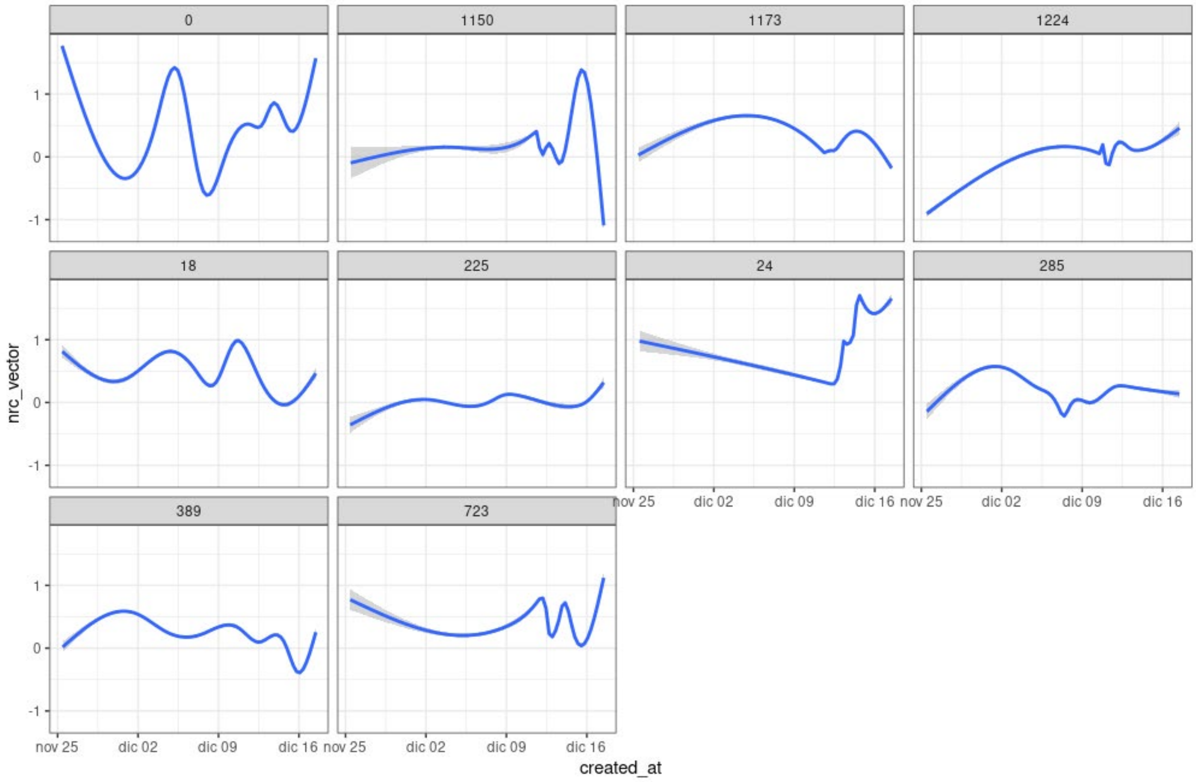

The emotional presence in the tweets is mainly one of trust and anticipation or rational thought, although disgust and fear also stand out. These basic emotions vary considerably depending on each group, as shown in Table 1. The @MichelleObama group stands for trust, followed far behind by @LeoDiCaprio, Greta Thunberg and anti-Thunberg followers. Anticipation is most prevalent in the group that tweets the Greta on Time Magazine's cover as Person of the Year. The highest levels of fear are found among the @MikeHudema environmental activists and the anti-Thunberg group, while disgust and sadness prevail mainly among the Spanish-language groups 225 and 285, which are at odds with each other.

Table 1: Presence of emotions by group. Source: Own elaboration

	0	1150	1173	1224	18	225	24	285	389	723
Anger	8.92%	9.52	4.90	10.39	8.12	14.36	4.29	15.09	12.03	14.93
Anticipation	11.71 %	22.04	41.45	13.53	13.94	7.65	5.81	14.74	15.09	13.42
Disgust	5.89%	4.17	3.02	8.94	4.47	23.03	2.13	9.10	6.69	7.07

Fear	14.25 %	12.31	7.60	14.56	16.66	12.67	4.37	14.77	16.08	10.64
Joy	9.61%	13.41	9.43	7.36	10.19	7.54	10.54	9.91	10.55	8.60
Sadness	9.00%	5.31	4.69	11.80	6.54	13.38	3.08	13.49	7.67	8.58
Surprise	8.00%	7.73	5.30	4.63	6.99	6.45	3.23	7.92	4.65	8.57
Trust	22.68 %	15.80	14.82	23.42	22.33	15.32	31.78	14.98	23.64	19.27

The longitudinal study of trust, disgust and fear shows that, except for the group around Greta Thunberg and the Time Magazine event and its consequences, the groups are more or less consistent. The groups around the US Democrats or the former first lady Michelle Obama experience an increase in trust after that date, while greater disgust is noted in the anti-Thunberg group. However, fear increases after the Time Magazine/ Donald Trump event in the pro-Thunberg group and similar ones.

The groups whose primary language is English see emotional movements and an increase in the most polarizing messages (those exceeding values of 3 in a positive or negative sense) after the publication of Time Magazine's Person of the Year. The chiefly Spanish- and Portuguese-speaking groups (225, 285, and 1224) slightly vary. By contrast, significant variations in emotions and polarity are found in the English-speaking groups, except the anti-Thunberg group, which increases its mean disgust value just days later. The movements among the Spanish groups occurred mainly before the publication of the Time magazine cover.

Hate Analysis

Hate was analyzed in English in all groups other than the Spanish groups (285 and 225). Figure 5 shows the measure obtained using a technique analogous to the previous point, where the mean level of hate in the discourse of the Spanish far-right group spirals upon Greta's arrival in Madrid and at the end of the Conference. Without attaining the same levels, the Democratic Party supporters (723), Michelle Obama supporters (24) and Greta Thunberg herself (0) see a rise in hate after Donald Trump's tweet on 11 December. Hate speech messages increased modestly in most of these groups towards the end of the Conference. Mean hate in the anti-Thunberg group in English (389) was low with little change over time.

Figure 5. Hate speech analysis from 11/25/2019 to 12/17/2019. Source: Own elaboration.

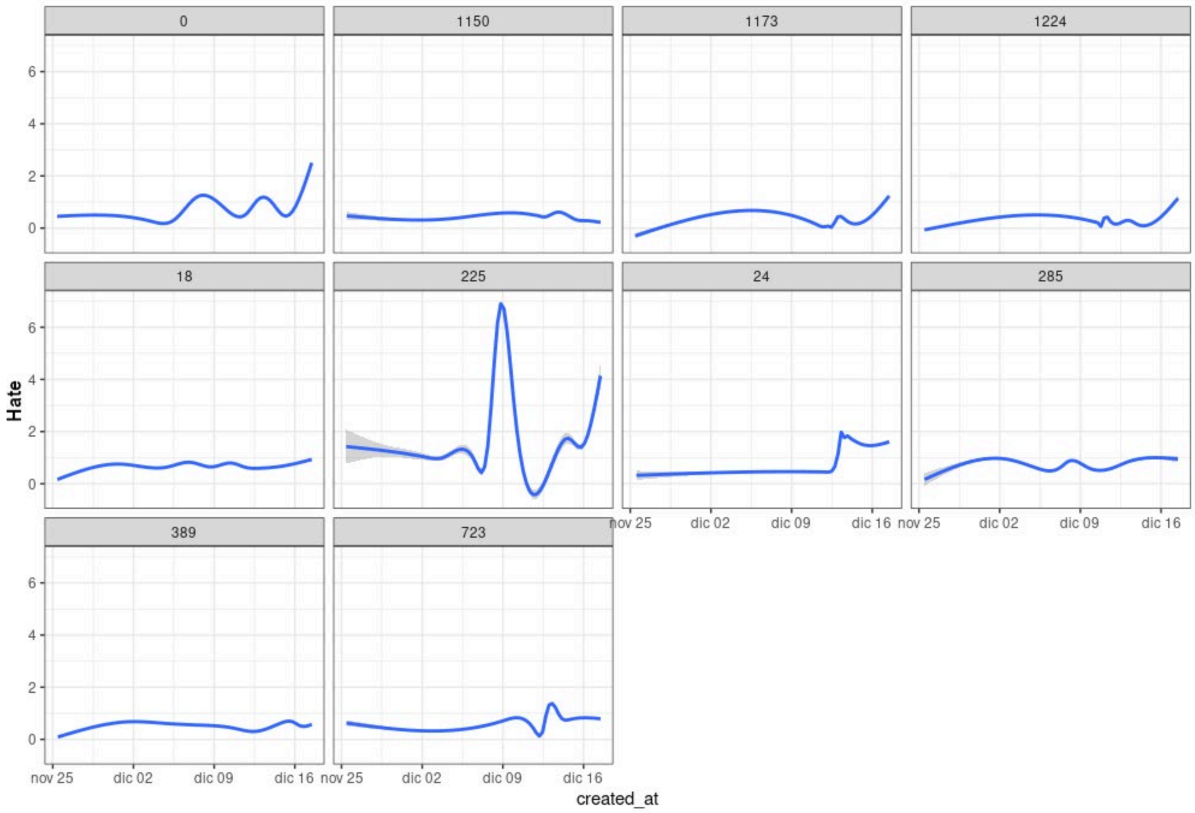

To determine which emotions were primarily responsible for hate speech in the posts, a tree-based regression model classified and numerically predicted the values obtained for this variable (Figure 6). The results illustrate the dependency of hate regarding feelings and polarity, in which the highest values of hate correspond to specific values of feelings linked to disgust (which determined the most intense levels of hatred), anger and sadness. The correlation between the prediction of the vectors and the true values, as well as the mean absolute error (MAE), determines whether the assignment adjusts correctly. The correlation of 0.6034 and an MAE of 0.7218 are good values.

Figure 6. Decision tree on total messages. The boxes below show the average hate intensity value, number of messages grouped together and their percentage over the total. Source: Own

elaboration

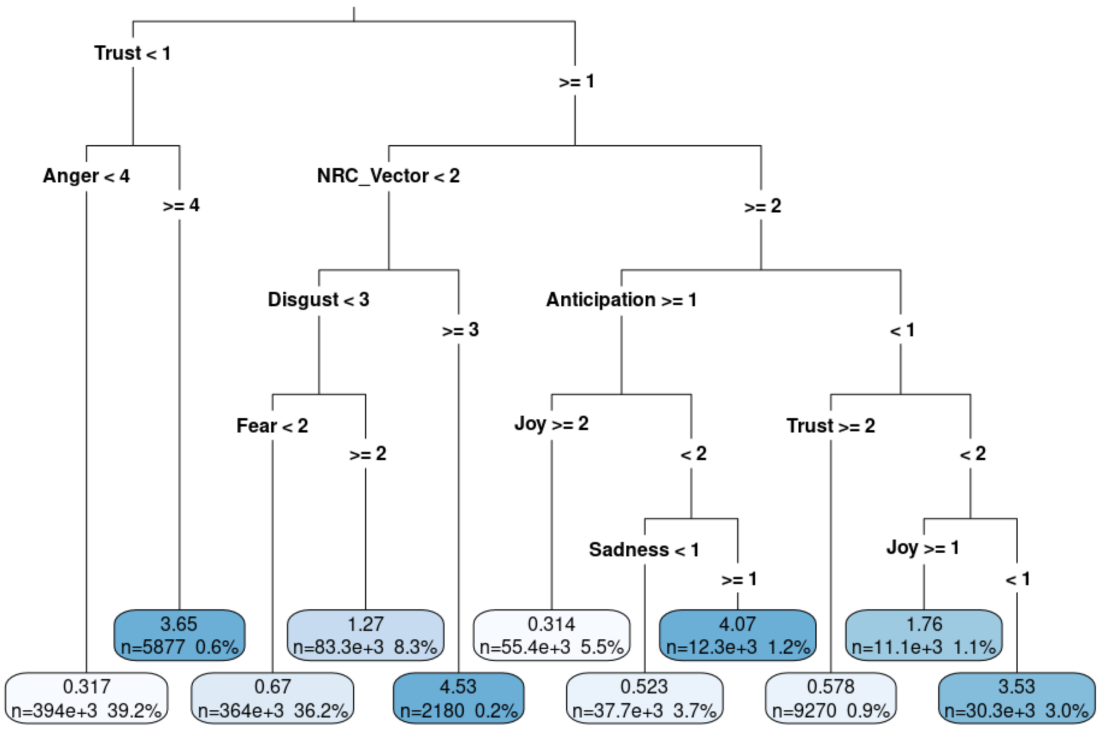

Overall, the study found that hate stems from disgust, anger, and sadness, but not in all the groups, nor to the same extent. Hatred does not arise from the same basic emotions, neither in quantity nor from the same message in each group. Thus:

- Group 0 (n=524,047): fear 18%, surprise 13%, sadness 13%, anticipation 12%, joy 12%, trust 11%, vector 10%, anger 5%, disgust 5%. Correlation: 0.6650, MAE: 0.5342.
- Group 723 (n=117,667): vector 30%, anger 23%, sadness 18%, fear 11%, trust 9%, surprise 4%, disgust 2%, joy 2%, anticipation 2%. Correlation: 0.57789, MAE: 0.5819.
- Group 1,173 (n=65,332): vector 28%, trust 18%, joy 16%, fear 8%, surprise 8%, anger 7%, anticipation 7%, sadness 6%, disgust 2%. Correlation: 0.6748, MAE: 0.4409.
- Group 389 (n=86,374): disgust 27%, vector 17%, anticipation 11%, trust, 10, anger 10%, fear 8%. joy 8%, sadness 4%, surprise 4%. Correlation: 0.2175, MAE: 0.7678.
- Group 24 (n=43,524): vector 28%, trust 22%, anger 10%, fear 8%, anticipation 8%, joy 8%, sadness 7%, surprise 6%, disgust 5%. Correlation: 0.5884, MAE: 0.5762.

- Group 18 (n=60,679): fear 24%, trust 16%, anger 12%, surprise 11%, vector 10%, sadness 9%, disgust 7%, anticipation 7%, joy 4%. Correlation: 0.4051, MAE: 0.6704.
- Group 1,224 (n=25,729): trust 25%, vector 18%, fear 15%, anger 12%, anticipation 10%, joy 9%, disgust 5%, sadness 3%, surprise 2%. Correlation 0.6658, MAE: 0.4891.
- Group 1,150 (n=39,743): disgust 24%, anger 21%, sadness 17%, trust 12%, fear 10%, vector 4%, anticipation 3%, joy 3%, surprise 2%. Correlation: 0.2049, MAE: 0.7028.
- Group 285 (n=27,039): trust 32%, fear 18%, disgust 15%, vector 11%, joy 8%, surprise 6%, anticipation 4%, anger 3%, sadness 3%. Correlation: 0.5260, MAE: 0.6578.
- Group 225 (n=15,471): trust 61%, fear 17%, vector 13%, joy 4%, surprise 2%, anticipation 2%. Correlation: 0.1116, MAE: 2.3764.

The values obtained by correlation or by MAE confirmed a moderate-good fit of the prediction made in each regression tree, with the weakest values in groups 1,150 and 225.

Group 225 shows that the highest level of hatred is mainly linked to trust and fear, while hate from the anti-Thunberg group (389) relates to disgust. By contrast, hate speech in the environmentalist groups (0 and 18) derives mainly from fear, while in the Democrat groups (723 and 24), it comes from polarity, trust and anger, and the rest from trust. Hence, we found that hate speech is not unique and does not stem from a single source but rather from a well-defined set of feelings, which vary according to the characteristics of each group and its opinions and interests.

No association was found with the variables determining the position of the message in the network, meaning that these motives do not statistically determine hate. Nevertheless, hate was shaped among the primary followers of the main accounts of the networks in each group, as shown in Figure 7. The accounts farthest away from the leading influencers (values below 0.25 in closeness centrality) of each group tweeted more messages with a greater intensity of hate. Those closest to them (values above 0.5) tweeted far less, except for the anti-Thunberg (389), Hong Kong (1150), and Spanish groups (225 and 285), the latter having followers who are more partisan and political.

Figure 7: Relationship between hate and centrality of proximity in the sender's network. Source: Own elaboration

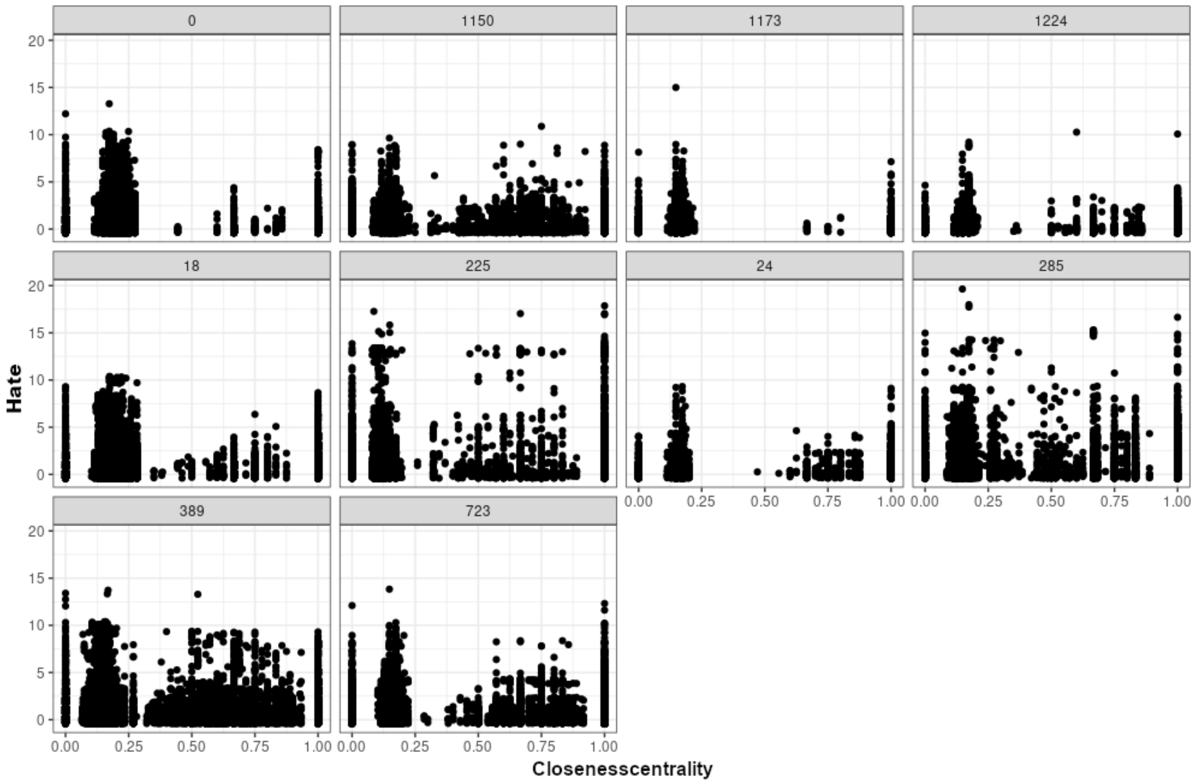

Study of accounts

The analysis of the accounts (Figure 8) shows that 26% of them have a 75% probability of being bots. The participating accounts were set up over several years, although many were created in 2009, late 2016, and late 2019. The algorithm analysis shows that almost a quarter of the accounts had a highly bot-like behaviour (whether human-controlled or automated).

Figure 8: Account creation date and histogram of bot behaviour. Source: Own elaboration

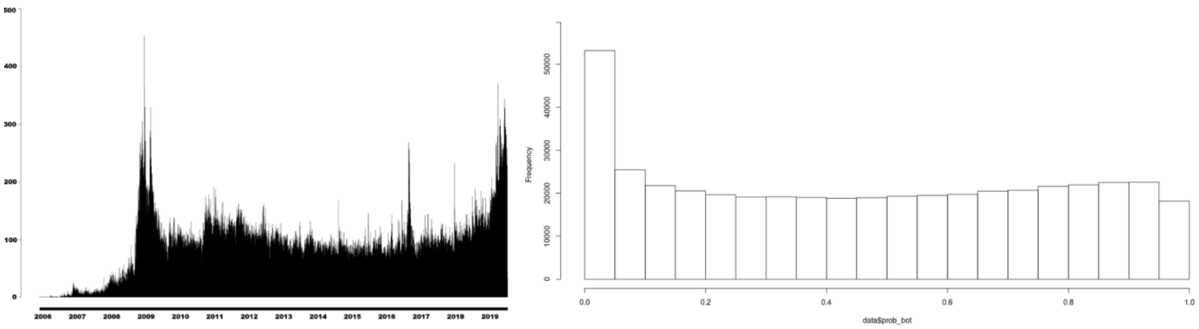

Conclusions

Our analysis of more than one million tweets on the young environmental activist Greta Thunberg's presence at the Climate Change Conference COP25 in Madrid, indicates a global impact. Notwithstanding, tweets referring to policies and leaders in specific countries show that the impact was greater in some places than in others. While the Conference was the principal focus of debate, US President Donald Trump's tweet taunting Thunberg as Time Magazine's Person of the Year 2019 marked a complete shift in the discourse, sparking a worldwide debate, though primarily in the United States, on climate change issues and the relationship between the two people. Accordingly, groups identified through Democratic leaders and former first lady Michelle Obama expressed their support for Thunberg. Those against her included members of the Republican Party, a Fox News columnist, and some accounts later deleted for not observing Twitter's rules. This spat saw an exponential multitude of tweets on November 11 and 12, diverting everyone's focus, except for Greta Thunberg and other environmentalists, away from the Conference.

This scenario corroborates previous studies (Berger and Milkman 2012; Weeks and Holbert 2013; Keib et al. 2018; Arce-García et al. 2020), namely the increased volume of tweets evoking emotions of greater intensity triggered users to retweet them. Further, the higher volume of tweets is also linked to changes in associated emotions in the messages, hate in particular, confirming that negative emotions can augment this effect, as noted in previous studies (Segado-Boj et al. 2020; Xue et al. 2020), in which feelings like fear or disgust act as catalysts to accelerate the spread of related content.

Debates with their own particularities also arose in specific countries, notably Spain, Brazil and Hong Kong. The debate in Spain was the most intense, especially in the far-right group (which had ties to the anti-Thunberg group in the United States), which showed the highest levels of hate, chiefly at the start of the Conference. In Brazil, the debate centred on President Bolsonaro and statements made by actor Leonardo DiCaprio, while in Hong Kong, the focus was on protests by local activists demanding change.

The discourse, measured from the intensity of its positive or negative valence and associated emotions, revealed a highly polarized debate, which intensified after Time Magazine announced Greta as its Person of the Year, and Trump's tweet in response to it. While trust was highly relevant in the principal leaders of each group, so was fear, as well as mutual disgust between the Spanish groups, and in the Democrat group after Trump's tweet.

Our study found a highly confrontational scenario, with its corresponding particularities in each area of the world in which it was debated, where the Greta Thunberg/ Donald Trump polemic took precedence over the climate change conference itself. This confirms the validity of the bubble filter (Pariser 2011) and echo chamber theories (Colleoni et al. 2014; Sunstein 2001; Jamieson and Cappella 2008) and how the confluence of the two leads to heightened polarization: well identified and separated groups according to network theory, and intense emotional values detected in the discourses, especially fear, disgust, trust and hatred among them. This finding derives from the relatively good degree of modularity for the Social Sciences (0.550) in assigning clusters, as well as the difference between the assignment of basic emotions and hatred in each one of the main groups identified, which increased in these groups at the end of the period under study. This ties in with the ideas of Arcila et al. (2017), according to which it is possible to anticipate on a greater scale how certain events will evolve.

Likewise, the regression tree models enabled us to predict the behaviour and chief emotions underlying the hate speech, which attained its highest levels when associated with anger, disgust, and sadness, albeit with evident particularities in each group. This technique of association by regression between hate and basic emotions is followed through the wheel of emotions derived by Plutchik (1980), and its study could open up new forms of analysis and deeper study to check what promotes hate in each group. Thus, hate expressed in the anti-Thunberg groups in the United States and in the Spanish ultra-right relates to disgust (the emotion generating the highest intensity of hate), trust and fear, whereas hate derived from fear among Thunbergs' environmentalist followers, and anger and trust among the American Democratic followers. Each group had its own reasons for expressing hate in their messages, although the predictive model shows that the most intense values were characterised by tweets associated with specific sentiments of an evident intensity.

We found that the accounts farthest from their influencers' network spread the most hate, in the form of individual discourse, apart from a few specific groups where hate is prevalent among close users with a strong connection to their network, in a more intense sentiment of following their bubble. The detection of hate spread by accounts far from their main influencers is explained by the fact that, although followers of their group, these accounts are close to a group network but hold own opinions; whereas the groups that were more partial to their cause provide a higher intensity and volume of hate speech to accounts closer to their influencers. The existence of many accounts created not long before the period under study or on specific dates, in addition to the high probability of bots, suggests an astroturfing phenomenon or orchestrated support campaign, or both, in some groups using co-tweeting and co-retweeting techniques (Keller et al. 2019) to evade detection by users, platforms and academic studies (Elmas et al. 2021).

To conclude, Donald Trump's tweet reproaching Greta Thunberg caused a strong emotional reaction and increased debate on Twitter. However, the reaction shifted the agenda from the debate on the climate emergency to inter-group confrontation, especially among the most partisan groups. This fulfils our objective of studying the incidence of hate speech and how it is structured: observing how it varies depending on each group and time, as a reaction to specific events that generate strong emotions.

The limitations of the study stem from the use of methods that are still under development, especially in non-English languages such as Spanish. These tools also tend to have certain problems with the use of irony or double meanings. In any case, the use of a large number of messages in the analysis means that the final result is not affected by any problem of language determination or double meaning in any of the texts. Future lines of work include the study of hate links and their dissemination through possible organised campaigns.

Bibliography

Arce-García S, Orviz-Martínez N and Cuervo-Carabel T (2020) Impacto de las emociones vertidas por diarios digitales en Twitter. *Profesional de la información* 29(5): e290520.

Arcila-Calderón C, Ortega-Mohedano F, Jiménez-Amores J and Trullenque S (2017) Análisis supervisado de sentimientos políticos en español: clasificación en tiempo real de tweets basada en aprendizaje automático. *El profesional de la información* 26(5): 973-982.

Agur C and Gan L (2021) Actors, Partisan Inclination, and Emotions: An Analysis of Government Shutdown News Stories Shared on Twitter. *Social Media+ Society* 7(2): 20563051211008816.

Alter C, Haynes S and Worland J (2020) Greta Thunberg, Person of the Year. The Power of Youth. *Time,* 11 December. Available at: https://time.com/person-of-the-year-2019-greta-thunberg/. Accessed 27 January 2022.

Barabasi AL (2016) *Network Science*. Cambridge (UK): Cambridge University Press.

Barrie C and Chun-ting HJ (2021) AcademictwitteR: an R package to access the Twitter Academic Research Product Track v2 API endpoint. *Journal of Open Source Software* 6(62): 3272.

Huth EJ, King K and Lock S (1988) Uniform requirements for manuscripts submitted to biomedical journals. *British Medical Journal* 296(4): 401–405.

Bassignana E, Basile V and Patti V (2018) Hurtlex: A multilingual Lexicon of Words to Hurt. In: *5th Italian Conference on Computational Linguistics. Collana dell'Associazione Italiana di Linguistica Computazionale*, 2253: 1-6. https://doi.org/10.4000/books.aaccademia.3085

Bastian M, Heymann S and Jacomy M (2009) Gephi: An Open Source Software for Exploring and Manipulating Networks. In: *Third International ICWSM Conference*. Available at: https://gephi.org/publications/gephi-bastian-feb09.pdf

Berger J and Milkman KL (2012) What makes online content viral? *Journal of Marketing Research* 49(2): 192-205.

Blondel V, Guillaume J, Lambiotte R and Lefebvre E (2008) Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008(10): 10008.

Bode L (2016) Political news in the news feed: Learning politics from social media. *Mass Communication and Society* 19(1): 24-48.

Breiman L, Friedman JH, Olshen RA and Stone CJ (1984) *Classification and Regression Trees*. Belmont: Wadsworth.

Chang WL (2019) The impact of emotion: A blended model to estimate influence on social media. *Information Systems Frontiers* 21: 1137-1151.

Chen J, Hossain S and Zhang H (2020) Analyzing the sentiment correlation between regular tweets and retweets. *Social Network Analysis and Mining* 10:13.

Colleoni E, Rozza A and 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.

Díaz-Campo J, Segado-Boj F and Fernández-Gómez E (2021) Hábitos del usuario y tipo de red social como predictores de consumo y difusión de noticias. *Profesional de la información* 30(4): e300417.

Duncombe C (2019) The politics of Twitter: emotions and the power of social media. *International Political Sociology* 13(4): 409-429.

Duffy A and Ling R (2020) The gift of news: Phatic news sharing on social media for social cohesion. *Journalism studies* 21(1): 72-87.

Ekman P (2003). Darwin, deception, and facial expression. *Annals of the New York Academy of Sciences*, 1000: 205-221.

Elmas T, Overdorf R, Özkalay AF and Aberer K (2021) Ephemeral Astroturfing Attacks: The Case of Fake Twitter Trends. *IEEE European Symposium on Security and Privacy*. Epub ahead of print 11 March. DOI: 10.1109/EuroSP51992.2021.00035.

Evolvi G (2017) #Islamexit: inter-group antagonism on Twitter. Information, *Communication & Society* 22(3): 386-401.

García-Perdomo V, Salaverría R, Kilgo DK and Harlow S (2018) To share or not to share: The influence of news values and topics on popular social media content in the United States, Brazil, and Argentina. *Journalism Studies* 19(8): 1180-1201.

Garland J and Chakraborti N (2012) Divided by a common concept? Assessing the implications of different conceptualizations of hate crime in the European Union. *European Journal of Criminology* 9(1): 38-51.

Gil de Zúñiga H, Weeks B and Ardèvol-Abreu Al (2017) Effects of the news-finds-me perception in communication: Social media use implications for news seeking and learning about politics. *Journal of Computer-Mediated Communication* 22(3): 105-123.

Hansen DL, Schneiderman B and Smith MA (2011) *Analyzing Social Media Networks with NodeXl. Insights from a Connected World.* Morgan Kaufmann: Massachusetts.

Jansson A and Lindell J (2015) News media consumption in the transmedia age. *Journalism studies* 16(1): 79-96.

Jacomy, M., Venturini, T., Heymann, S. & Bastian, M. (2014) ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software. *PLOS ONE* 9(6): e98679.

Jamieson KH and Cappella JN (2008) Balkanization of Knowledge and Interpretation. In: Jamieson KH and Cappella JN (eds) *Echo Chamber: Rush Limbaugh and the conservative media establishment*. New York: Oxford University Press.

Jockers M (2017) Syuzhet, extracts sentiment and sentiment-derived plot arcs from text. Available at: https://rdrr.io/cran/syuzhet/. Accessed 26 January 2022.

Karlsen R (2015) Followers are opinion leaders: The role of people in the flow of political communication on and beyond social networking sites. *European Journal of Communication* 30(3): 301-318.

Kearney MW (2018) Tweetbotornot: An R package for classifying Twitter accounts as bot or not. Available at: https://github.com/mkearney/tweetbotornot. Accessed 27 January 2022.

Kearney MW (2019) Rtweet: Collecting and analyzing Twitter data. *Journal of Open Source Software* 4(42): 1829.

Keib K, Himelboim I and Han JY (2018) Important tweets matter: Predicting retweets in the #BlackLives-Matter talk on Twitter. *Computers in human behavior* 85:106-115.

Keller F, Schoch D, Stier S and Yang J (2019) Political Astroturfing on Twitter: How to Coordinate a Disinformation Campaign. *Political Communication* 1 (25): 256-280.

Lantz B (2019) *Machine Learning with R.* Birmingham: Packt Publishing.

Mohammad S and Turney P (2010) Emotions Evoked by Common Words and Phrases: Using Mechanical Turk to Create an Emotion Lexicon. In: Inkpen D and Strapparava C (eds) *Proceedings of the NAACL-HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text,* pp.26-34.

Mohammad S and Turney P (2013) Crowdsourcing a Word-Emotion Association Lexicon. *Computational Intelligence* 29(3): 436-465.

Noor P (2019). Trump's latest attack on Greta Thunberg was sexist, ableist – and perhaps jealous. *The Guardian,* 12 December. Available at: https://www.theguardian.com/environment/2019/dec/12/donald-trump-greta-thunberg-time-magazine. Accessed 27 January 2022.

Oltmann SM, Cooper TB and Proferes N (2020) How Twitter's affordances empower dissent and information dissemination: An exploratory study of the rogue and alt government agency Twitter accounts. *Government Information Quarterly* 37(3): 101475.

Pariser E (2011) *The filter bubble: How the new personalized web is changing what we read and how we think*. New York: Penguin.

Paz MA, Montero-Díaz J and Moreno-Delgado A (2020) Hate speech: A systematized review. *Sage Open* 10(4): 2158244020973022.

Picone I, Courtois C and Paulussen S (2015) When news is everywhere. Understanding participation, cross-mediality and mobility in journalism for a radical user perspective. *Journalism Practice* 9(1): 35-49.

Plutchik R (1980). A general psychoevolutionary theory of emotion. *Emotion: Theory, Research, and Experience*, 1(3): 3–33. https://doi.org/10.1016/B978-0-12-558701-3.50007-7

Reuters (2020). Climate activist Thunberg hits back at Trump over anger management taunt. *Reuters,* 6 November. Available at: https://www.reuters.com/article/us-usa-election-thunberg-idUSKBN27M0TN. Accessed 27 January 2022.

Sauter DA, Eisner F, Ekman P and Scott S (2010) Cross-cultural recognition of basic emotions through nonverbal emotional vocalizations. *Proceedings of the National Academy of Sciences* 107(6): 2408–2412.

Rodrigues JPC (2021) *Advances in Delay-Tolerant Networks (DTNs): Architecture and enhanced performance.* Cambridge: Elsevier.

Segado-Boj F, Díaz-Campo J, Navarro-Asensio E and Remacha-González L (2020) Influence of News-Finds-Me Perception on accuracy, factuality and relevance assessment. Case study of news item on climate change. *Revista Mediterránea de Comunicación* 11(2): 85-103.

Serrano-Puche J, Fernández CB and Rodríguez-Virgili J (2020) Información política y exposición incidental en las redes sociales: un análisis de Argentina, Chile, España y México. *Doxa comunicación. Revista interdisciplinar de estudios de comunicación y ciencias sociales* 27: 19-42.

Sunstein CR (2001*) Republic 2.0.* Princeton: Princeton University Press.

Sunstein CR (2017) *#Republic. Divided Democracy in the Age of Social Media.* Princeton University Press.

Swati U, Pranali C and Pragati S (2015) Sentiment analysis of news articles using machine learning approach. *International Journal of Advances in Electronics and Computer Science 2*(4): 114-116.

Tausczik YR and Pennebaker JW (2010) The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology* 29(1): 24-54.

Urman A (2019) Context matters: political polarization on Twitter from a comparative perspective. *Media, Culture & Society* 42(6): 857-879.

Wasike B (2017) Persuasion in 140 characters: Testing issue framing, persuasion and credibility via Twitter and online news articles in the gun control debate. *Computers in Human Behavior* 66: 179-190.

Weeks BE and Holbert RL (2013) Predicting dissemination of news content in social media: A focus on reception, friending, and partisanship. *Journalism & Mass Communication Quarterly* 90(2), 212-232.

Williams M (2021) The Science of Hate. Faber & Faber: London. ISBN: 978 0 571 35706 2

Xue J, Chen J, Hu R, Chen C, Zheng C, Su Y and Zhu T (2020) Twitter discussions and emotions about the COVID-19 pandemic: Machine learning approach. *Journal of Medical Internet Research* 22(11): e20550.


DOUBLE ISSUE
DEC. 23 / DEC. 30, 2019
PERSON of the YEAR
TIME
GRETA
THUNBERG
THE POWER
OF YOUTH
time.com