Uphill from here: Sentiment patterns in videos from left- and right-wing YouTube news channels

Felix Soldner[1], Justin Chun-ting Ho[2], Mykola Makhortykh[3],
Isabelle van der Vegt[1], Maximilian Mozes[1,4] and Bennett Kleinberg[1,3]
[1]University College London
[2]University of Edinburgh
[3]University of Amsterdam
[4]Technical University of Munich
[1]{felix.soldner, isabelle.vandervegt, bennett.kleinberg}@ucl.ac.uk
[2]justin.ho@ed.ac.uk, [3]m.makhortykh@uva.nl, [4]mozes@cs.tum.edu

Abstract

News consumption exhibits an increasing shift towards online sources, which bring platforms such as YouTube more into focus. Thus, the distribution of politically loaded news is easier, receives more attention, but also raises the concern of forming isolated ideological communities. Understanding how such news is communicated and received is becoming increasingly important. To expand our understanding in this domain, we apply a linguistic temporal trajectory analysis to analyze sentiment patterns in English-language videos from news channels on YouTube. We examine transcripts from videos distributed through eight channels with pro-left and pro-right political leanings. Using unsupervised clustering, we identify seven different sentiment patterns in the transcripts. We found that the use of two sentiment patterns differed significantly depending on political leaning. Furthermore, we used predictive models to examine how different sentiment patterns relate to video popularity and if they differ depending on the channel's political leaning. No clear relations between sentiment patterns and popularity were found. However, results indicate, that videos from pro-right news channels are more popular and that a negative sentiment further increases that popularity, when sentiments are averaged for each video.

Keywords: linguistic temporal trajectory analysis, online news, left-wing, right-wing, sentiment analysis, YouTube

1 Introduction

Today, news is increasingly consumed through online platforms. Approximately nine out of ten adults (93%) in the US acknowledge reading news online (Center, 2018); similar, if slightly lower rates, are observed for many European countries (e.g. 65% of Germany and 79% in the Netherlands) (Newman et al., 2018). The reasons for the growing online consumption of news are many: digital outlets provide higher interactivity and greater freedom of choice for news readers (Van Aelst et al., 2017), and the possibility to access them through mobile devices enables more flexibility and convenience in consuming news content (Schrøder, 2015).

Despite the multiple benefits associated with "the digital turn" in news consumption, it also raises concerns regarding possible changes in the societal role of news media (Coddington, 2015). It has been established that there is a relationship between news consumption and the formation of public agendas, including political attitudes of the general population (McCombs, 2014). However, the possible impact of digitalization on these relationships remains a subject of academic debate (Helberger, 2015; Eskens et al., 2017; Van Aelst et al., 2017). An increasing number of studies examine the possible connection between online news consumption and the formation of isolated ideological communities, which can nurture biases and limit citizens' societal participation (Gentzkow and Shapiro, 2011; Flaxman et al., 2016; Sunstein, 2017). Such formations can lead to increased political or personal radicalization, which is a concern for societal cohesion.

Polarization concerns related to audience segmentation are amplified by the use of affective language for producing online news stories (Soroka et al., 2015). Despite the widespread belief that news stories tend to use balanced and unbiased language, recent studies suggest that this is not necessarily true, in particular in the case of news stories crafted for online environments such as news websites (Young and Soroka, 2012). A number of studies suggest that the use of affective language can increase audience engagement with

84

news content (e.g. stories with a negative tone seem to be more popular (Trussler and Soroka, 2014; Soroka et al., 2015)). At the same time, the relationship between a narrative's popularity and the specific patterns of sentiment remains under-investigated, as well as the ways this relationship varies among audiences with different political leanings.

In this paper, we contribute to the debate on the use of emotionally loaded language in online news stories as well as its variations between outlets with different political leanings. Specifically, we examine the sentiment trajectories of YouTube news channels related to major English-language media outlets. Our interest in YouTube is attributed to the platform's significant impact on information consumption as users increasingly consume news in the video format (al Nashmi et al., 2017). Together with Facebook, YouTube is the main platform used by consumers to watch news outside of news organizations' own websites (Newman et al., 2018). News content produced by news organizations for YouTube usually follows traditional broadcast standards (Peer and Ksiazek, 2011) and formats (al Nashmi et al., 2017) and often serves as an extension of their distribution model. At the same, YouTube also enables experimenting with new formats and engaging audiences in more interactive or provocative ways (al Nashmi et al., 2017).

Using YouTube as a case platform, we test for differences in the dynamic use of sentiment in news coverage between channels leaning to the political left (e.g. CNN) and right (e.g. Fox News). Additionally, we examine if sentiment patterns have predictive value for news videos' popularity on YouTube.

2 Related Work

Since the mid 2000s, sentiment analysis remains one of the fastest growing fields of machine learning and computational linguistics (Mäntylä et al., 2018). Defined by Liu (2010) as a collection of methods for detecting and extracting subjective information (e.g. opinions, attitudes and emotions) from language, sentiment analysis is increasingly adopted for multiple academic areas, varying from from media studies (Sivek, 2018) to literature (Gao et al., 2016) to political science (Cambria, 2016) and conflict studies (Welch, 2018).

The growing adoption of sentiment analysis as a research tool is accompanied by methodological improvements. Originally employed for classifying user reviews coming from the commercial domain (Pang and Lee, 2008), early approaches to sentiment analysis were focused on producing a single sentiment score for the specific document or a text section using binary assessments of polarity. Such approaches, however, resulted in rather simplified evaluations, which reduced sentiment complexity to binary constructs under which the whole document could be either positive or negative or (in some cases) neutral. This can lead to incorrect conflations of sentiment variations within a text.

The response to these limitations included the advancement towards more fine-grained assessments of sentiment, including the identification of a more complex spectrum of emotions (Cambria et al., 2015), but also the transition towards the dynamic assessment of sentiment shifts throughout texts (Jockers, 2015; Tanveer et al., 2018; Kleinberg et al., 2018).

Until now, the sentiment of news stories remains a rather under-investigated subject. Kaya et al. (2012) note that unlike user reviews of products (e.g. movies), news stories are considered to be written in a neutral way. A different study found that news stories about the economy and the environment were overall more positive than about crime or international topics, which were overall more negative (Young and Soroka, 2012). Despite the growing number of studies on news sentiment, only a few of them so far approach it by considering the dynamic shifts of sentiment. A study, examining sentiments of different topics (e.g. earthquakes) included time stamps over a three months period, to model the sentimental change of news stories (Fukuhara et al., 2007). In that way, they were able to show a sudden increase of negative emotions in news corpora, when an earthquake occurred. These negative emotions slowly decreases over time. While this approach models sentiment change over time, it cannot account for shifts within single stories.

Among multiple areas of research on news story sentiment, the issue of variation in the use of affective language by news outlets with different political leanings (e.g. left- or right-wing) is both under-investigated and urgent. The language of politics-related texts does not only reflect, but can also influence the sentiments of the audience (Brader,

2005). A number of studies point to the distinct features of online political communication depending on the actors' political leaning (e.g. (Engesser et al., 2017; Bracciale and Martella, 2017; Hameleers et al., 2017; Schoonvelde et al., 2019a)). Engesser et al. (2017) identify that social media usage of right-wing political actors is characterized by the use of a few key features, such as emphasizing the sovereignty of the people, advocating for the people, attacking the elite, ostracizing others, and invoking the concept of a "heartland". Similarly, Bracciale and Martella (2017) show that communication styles of populist actors tend to involve highly emotional language and that they are particularly keen on referring to negative emotions (e.g. fear) to mobilize their supporters.

The studies mentioned above, however, focus on political statements produced and distributed through different media; yet, little is known how different political sentiments materialize in online news stories, in particular in the YouTube video format, which is the major focus of our study. A number of studies discuss the impact of news outlet ideological leanings on the way specific subjects are covered (de Vreese, 2005). Most of these works, however, tend to focus on traditional formats of news stories (i.e. text) and use qualitative approaches to examine coverage of a specific subject, such as climate change (Dotson et al., 2012; Feldman et al., 2012) or protest campaigns (Ha and Shin, 2016; Shahin et al., 2016). In our paper, we propose to look at the intra-textual dynamics of the overall sentiment of content produced by the outlet in question and employ a quantitative approach to trace if there are differences between right- and left-wing news outlets.

3 Method

The data used in this study are publicly available[1] and the current work is the joint product of a workshop on linguistic temporal trajectory analysis at the European Symposium Series on Societal Challenges in Computational Social Science in 2018. This includes the pre-processing and feature extraction[2] of the data, which is needed for the analyses we performed in the current study. The data has not been used in other research and was specifically collected to devise the current work.

3.1 News channels selection

The data consist of all English-language channels from the top 250 news channels on YouTube, which were ranked by SOCIAL-BLADE[3] (`www.socialblade.com`, retrieved November 2018). From that pool, 18 news channels were selected, which were identified as the ones holding political bias (i.e. either left- or right-wing) by Media Bias/Fact Check[4]. The website uses a rating method to identify biases among information sources and has been used by previous studies dealing with media bias (Bentley et al., 2019; Bovet and Makse, 2019; Mehta and Guzmán, 2018).

3.2 Obtaining video transcripts

Video transcripts were scraped with the help of "`www.downsub.com`", which retrieves the transcripts of specific YouTube video URLs, and the "beautifulsoup" python package (Richardson, 2019) (for more details see Kleinberg et al. (2018)). Downloaded transcripts were manually or automatically generated. Videos without transcripts were not included in further analyses. Retrieved transcripts were cleaned by removing XML tags and merged into one string, without punctuation for each video. Selected transcripts for further processing adhered to the following criteria: at least 100 words; at least 50% of words are matched English words; at least 90% of words are ASCII-encoded; and from channels with more than 2000 valid transcripts. Subsequently, for 4 left and 4 right channels (randomly selected) 2000 transcripts were selected, 7 non-English transcripts from Business Insider and Russian today were then excluded, resulting in a balanced dataset of 15993 transcripts (table 1).

3.3 Popularity rating

Since we are interested in the popularity rating of each video and its association to sentiment style, we created an adjusted popularity rating. We calculated a popularity index defined as the number of upvotes divided by the total views for each video, which would be a score between 0 (no upvotes) and 1 (upvotes is equal to views). The adjustment allows us to compare the popularity be-

---

tween videos, which have different upload dates, because videos available over a prolonged period of time have the advantage of accumulating more upvotes.

3.4 Feature extraction

The organizers of the workshop provided us with the necessary features for further analyses, who were inspired by Jockers (2015) and Gao et al. (2016); (for more details see Kleinberg et al. (2018)). Features were generated from the transcripts, which capture the sentiment change throughout the transcript. The applied method is based on the approach of the R package "sentimentr" (Rinker, 2019a), which generates sentiments on a sentence level, but the current approach extends it to continuous text without punctuation as is the case with video transcripts. The "naive context" sentiment extractor (Kleinberg et al., 2018) accounts for valence shifters, which influence the meaning of the sentiment. Negators (e.g., not, doesn't), [de-]amplifiers (e.g., really, hardly), and adversative conjunctions (e.g., but, however) were included. This is important when generating the sentiments for sentences like "I had a really good day" (amplifying "good" through "really") or "My day was not bad" (changing "bad" from a negative to a positive sentiment through "not"). In order to extract sentiment values, each word is assigned a sentiment value based on the "Jockers and Rinker Polarity Lookup Table" from the lexicon R package (Rinker, 2019b). For the extractions of the features a window approach is taken, by which the sentiment of the core word and its surrounding words of $+/-3$ are considered. This cluster of seven words is assigned an adjusted sentiment value by calculating the product of all the sentiment values in the cluster. The features are represented in vectors, containing zeros and weighted sentiment values. Theses values are standardized to a narrative time from 0 to 100, with a discrete cosine transformation from the "syuzhet" R package (Jockers, 2015) and scaled from -1 to $+1$ (i.e. lowest sentiment (negative) per transcript to highest sentiment (positive) per transcript).

4 Results

4.1 Data

The total data set consists of 15993 video transcripts, which are distributed between eight news channels, and are equally divided into political

Channel	Pol	N. of videos	Avg. wc
Al Jazeera English	Left	2000	1000.49
Business Insider	Left	1997	501.56
Fox news channel	Right	2000	777.85
MSNBC clean forward	Left	2000	1499.25
Russia today[5]	Right	1996	1422.71
The young turks	Left	2000	1504.59
The daily wire	Right	2000	4845.80
Rebel Media	Right	2000	1250.73

Table 1: YouTube news channels distribution; Pol = Political leaning; N. = number; wc = word count

leanings (table 1).

4.2 Clustering

In order to examine sentiment patterns within the YouTube videos and their predictive value on popularity, we examined whether overarching sentiment patterns are present. We used an unsupervised k-means clustering method and determined the number of k through the within cluster sum of squares inflexion method (Thorndike, 1953) for 1 to 30 clusters (figure 1). Following this, we used a k-means method with $k = 7$, which resulted in seven clusters, each representing a sentiment behavior pattern found in the transcripts. All the patterns are displayed in figure 1, showing the average sentiment trajectory for the adjusted timescale from 0 to 100, and sentiments from -1 (most negative) to $+1$ (most positive). The patterns correspond well with the patterns found in related work (Kleinberg et al., 2018). Thus, we assigned the corresponding taxonomy names to our patterns. The dotted blue line indicates the average behavior pattern for the cluster and the red lines indicate the standard deviation of $+/-1$. Table 2 shows the taxonomies and descriptive statistics of the sentiment clusters.

4.3 Sentiment styles and political stance

We also examined whether there is a relationship between political stance and sentiment clusters. A significant association was found with a 2 (political stance) by seven (cluster) Chi-square test ($\chi^2(14) = 25.31, p < 0.001$). Table 3 shows that the cluster "Downhill from here" was significantly $p < 0.01$ more used by politically right leaning news channels. The reversed effect was observed for the cluster "Uphill from here", which

---

[5]Russia today (RT) can be considered as having a right-leaning political political stance at its core. However, this can have some exceptions, such as supporting the yellow vest movement in France.
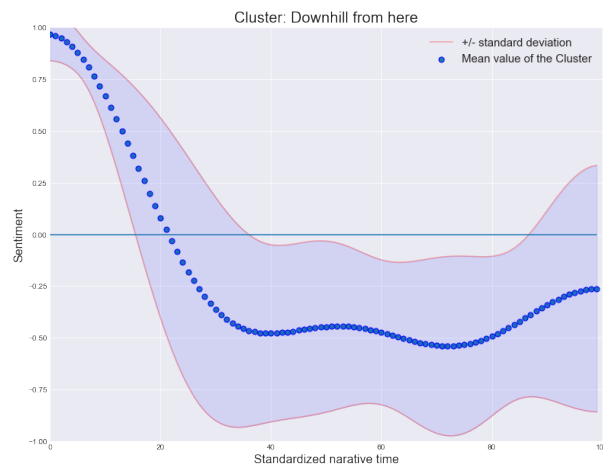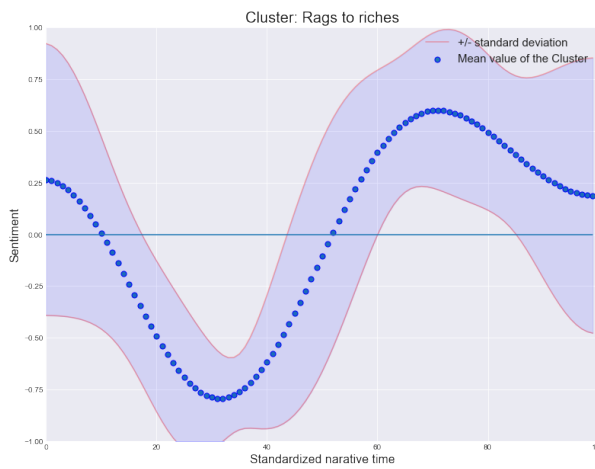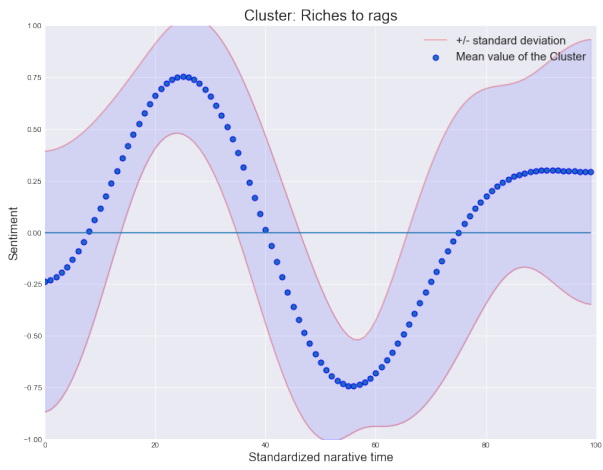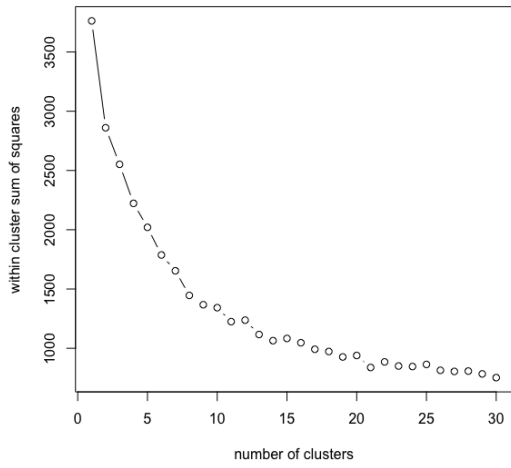
Figure 1: Cluster plot and sentiment behavior patterns for each cluster.

Cluster	Description	N. of videos	% of videos	Avg. vc	Avg. up v.
Rags to riches	Negative curve turns into positive curve	2675	16.73	827.00	15.53
Riches to rags	Positive curve turns into negative curve	2587	16.18	1002.47	17.11
Downhill from here	Short positive turns into consistent negative	2177	13.61	919.94	16.82
End on a high note	Short negative turns into consistent positive	2194	13.72	928.78	17.03
Uphill from here	Consistent negative turns into short positive	2085	13.04	823.17	14.37
End on a low note	Consistent positive turns into short negative	1547	9.67	846.80	16.81
Mood swing	Small positive start into negative-positive-negative curves with small positive ending	2728	17.06	910.83	16.57
Total	All	15993	100	897.71	16.32

Table 2: Sentiment styles taxonomy (adopted from Kleinberg et al. (2018)) and descriptive statistics; Average (Avg.) scores are adjusted by the number of days the videos were uploaded; N = Number; vc = view count; v = votes.

	Political leaning	
Cluster	Left	Right
Rags to riches	-0.96	0.96
Riches to rags	1.09	-1.09
Downhill from here	-3.25*	3.25*
End on a high note	1.28	-1.28
Uphill from here	2.74*	-2.74*
End on a low note	-2.44	2.44
Mood swing	1.13	-1.13

Table 3: Chi-Square residuals; * = statistically significant ($\alpha = 0.01$).

were used more often by politically left leaning news channels.

4.4 Sentiment clusters and popularity

To assess the relationship between the sentiment clusters and political orientations on popularity rating, we conducted three least square regression models in R with the "caret" package (Kuhn, 2008). The different sentiment clusters were the predictors for the adjusted popularity rating. We used the cluster "mood swing" as the reference category as it was closest to the overall average of the adjusted upvotes, all other clusters were treated as a separate dichotomous variable. Our first regression model included sentiment clusters, which consisted of all news channels. The second regression model included the channel as a fixed effect along with sentiment clusters. The analysis indicated that there was no significant difference in the adjusted popularity rating between the clusters in the second model; "Rags to riches" ($\beta = -1.28, se = 1.60, p = .42$), "Riches to rags" ($\beta = .47, se = 1.61, p = .77$), "Down hill from here" ($\beta = -1.25, se = 1.69, p = .46$), "End on a high note" ($\beta = .4, se = 1.69, p = .81$), "Up hill from her" ($\beta = -1.71, se = 1.71, p = .32$), "End on

a low note" ($\beta = .02, se = 1.87, p = .99$). Neither model explains a sufficient proportion of the variance to be considered informative, $R^2 = 0.00$ and $R^2 = 0.03$, for the first and second model, respectively.

We also split the transcripts in three equal sized components (beginning, middle, and end) and calculated the average sentiment rating for each part and used a OLS regression model to test for an effect of the components on the adjusted popularity ($F(10, 15982) = 52.46, p < .001, r^2 = 0.03$). No significant effects were found, after controlling for channel: "Beginning" ($\beta = -1.28, se = 1.02, p = .68$), "Middle" ($\beta = -1.65, se = -1.65, p = .11$), "End" ($\beta = -1.79, se = 1.09, p = .1$).

In addition we used an OLS regression model to test if the average sentiment score of each transcript and political leaning had an effect on the adjusted popularity ($F(3, 15989) = 1759, p < .001, r^2 = 0.248$). It seems that the model can account for 24% of the variance in adjusted popularity. The model exhibits a significant constant ($\beta = 0.013, se < 0.001, p < 0.001$), a significant main effect of political leaning (right) ($\beta = 0.021, se < 0.001, p < 0.001$), an insignificant main effect for average sentiment ($\beta = -0.0004, se = 0.001, p = 0.66$), and a significant interaction between average sentiment and political leaning (right) ($\beta = -0.003, se < 0.001, p = 0.034$). The model predicted adjusted popularity for left wing channel when average sentiment equals to zero is 0.013 and 0.013 + 0.02 = 0.033 for right wing channels. The slope of the regression line for the left wing channel is -0.0004 and -0.0004 - 0.003 = -0.0034 for right wing channels, suggesting that the effect of average sentiment is greater in magnitude for right wing than for left wing channels.

5 Discussion

In this study we examined sentiment patterns in news videos published on YouTube channels with different political leanings. Using sentiment trajectory analysis, we identified recurring patterns of sentiment changes, which were then grouped through *k*-means clustering into seven major categories of videos based on their sentiment patterns (e.g. "rags to riches" or "mood swings"). These patterns correspond to the ones identified in previous research on sentiment patterns of vloggers on YouTube (Kleinberg et al., 2018). While YouTube videos of vloggers and news channels differ in their domain, the persistence of similar sentiment clusters might indicate the presence of a few consistent sentiment styles that are shared between specific content domains and themes across YouTube. Future research could further examine whether similar patterns persist across various domains of YouTube content.

5.1 Sentiment pattern by political stance

Our results show, that political leanings seem to influence the usage of sentiment patterns: "Downhill from here" was used more often by pro-right news channels than by pro-left news channels, and conversely "Uphill from here" was used more often by pro-left news channels. It is interesting that both sentiment patterns exhibit the same proportion of negative and positive sentiment (80/20 respectively), but are different in sentiment order (see figure 1).

It is important to note that previous studies show that the sentiment of politics-related content, such as political ads, can affect the emotional state of the viewer (Brader, 2005). For the "Downhill from here" and "Uphill from here" patterns, viewers of pro-right news channels are left with a more negative sentiment and viewers of pro-left news channels are left with a more positive sentiment after watching the whole video.

At the same time, while some sentiment patterns seem to be more commonly used by channels with specific political leaning, it is difficult to propose a strong theoretical background which can explain this link. The complexity of this task is related to the large number of factors (both ideological, but also contextual) which can influence the use of language for political purposes (for a more detailed discussion see recent work on linguistic complexity of political speeches (Schoonvelde

et al., 2019b)).

Further research could examine the distribution of sentiment patterns between specific YouTube channels and investigate how factors, such as viewership or content influence them. This could include how pro-right/left channels differ in the amount of content creation for specific topics and how sentiment is used within this content. In addition, future work could examine if political leaning has predictive value on sentiment clusters for specific topics. This could be useful to ascertain whether news channels with a right or left political leaning discuss specific topics more often and differently than others and with what type of sentiment style.

5.2 Predicting video popularity through sentiment patterns

Our study indicated that sentiment patterns are weak predictors of news video popularity. A number of studies suggest that content-based features, in particular sentiment, have a strong impact on news content popularity (Trussler and Soroka, 2014; Soroka et al., 2015). However, our findings align with results of other studies that emphasize the importance of looking at a broader set of features, in particular contextual ones (e.g., the time of publication), for predicting the popularity of news content (Tatar et al., 2012; Keneshloo et al., 2016). Additionally, in the case of YouTube, the importance of other content-agnostic factors (e.g., the total views a channel received previously) for predicting the popularity of videos has been noted (Borghol et al., 2012; Figueiredo et al., 2014).

It is important to note that the current prediction task, with sentiment clusters, does not account for the temporal properties of the sentiment trajectories. Future research could utilize the sequential alignment of the raw sentiment scores by integrating this aspect into a prediction task of video popularity. That way, the features could be used without condensing information (e.g., into clusters) and acknowledge the sequential nature of sentiment trajectories.

5.3 Predicting popularity through average sentiment and political stance

We also tested for effects of average sentiment scores and political leaning on popularity with. The regression model was able to account for 24% of the variance of video popularity. Examining the model's coefficients more closely show, that

popularity seem to be higher when the video originated from a pro-right news channel. Furthermore, the interaction of political leaning and sentiment scores shows, that a negative sentiment will increase popularity, while a positive sentiment will decrease popularity more for videos from pro-right news channels. Our findings support earlier work, which show that a negative tone seem to be more popular overall (Trussler and Soroka, 2014; Soroka et al., 2015).

5.4 Limitations

The current dataset consists of transcripts of YouTube videos and it is important to recognize that aspects such as video and audio of the clips are not integrated in the analyses. News channels might utilize audio and visual effects differently, which could affect text sentiment and video popularity.

In addition, the obtained transcripts in our analyses could have been generated manually or automatically, hence might differ in quality. Since there is no direct indicator of this, we do not know in what proportion they are represented in our corpus.

Generating appropriate and accurate bias ratings for news channels is not easy. Therefore, it is not guaranteed that the bias rating of Media Bias/Fact Check is accurate in all regards. However, it has a comprehensive list of news channels, which are not always covered by other bias rating resources (Budak et al., 2016; Center, 2018).

Finally, in our study we specifically focused on YouTube videos. While earlier studies (Peer and Ksiazek, 2011; al Nashmi et al., 2017) demonstrate that content produced by legacy media for YouTube often follows the same standards and formats as stories produced for other platforms, there also exceptions from this rule. Some news organizations (for instance, RT) tend to push more provocative stories to YouTube, whereas others (such as CNN) preferred to publish more lighter content on the platform (al Nashmi et al., 2017). These distinct features of content distributed through YouTube news channels can impact our observations.

6   Conclusion

In this study we showed that news channels on YouTube exhibit different sentiment patterns, which can be clustered into overarching groups.

We found seven sentiment shapes similar to those found in previous research. The cluster "Mood swings" was most prominent whereas "End on a low note" was least prominent. Two additional sentiment clusters seemed to be used differently depending on political leaning of the channels: the cluster "Downhill from here" was used more often by pro-right news channels than by pro-left news channels. The reversed effect was observed for the cluster "Uphill from here". In addition, sentiment clusters seem to have no predictive value on popularity ratings. However, we found that pro-right videos were more popular and that negative sentiments increased popularity, for averaged sentiment scores of each video. Future research on dynamic approaches to sentiment analysis might help overcome some of the current limitations and offer more nuanced insights into language use in online media.

References

Frank Bentley, Katie Quehl, Jordan Wirfs-Brock, and Melissa Bica. 2019. Understanding Online News Behaviors. page 11.

Youmna Borghol, Sebastien Ardon, Niklas Carlsson, Derek Eager, and Anirban Mahanti. 2012. The untold story of the clones: Content-agnostic factors that impact YouTube video popularity. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '12*, page 1186, Beijing, China. ACM Press.

Alexandre Bovet and Hernán A. Makse. 2019. Influence of fake news in Twitter during the 2016 US presidential election. *Nature Communications*, 10(1).

Roberta Bracciale and Antonio Martella. 2017. Define the populist political communication style: The case of Italian political leaders on Twitter. *Information, Communication & Society*, 20(9):1310–1329.

Ted Brader. 2005. Striking a Responsive Chord: How Political Ads Motivate and Persuade Voters by Appealing to Emotions. *American Journal of Political Science*, 49(2):388.

Ceren Budak, Sharad Goel, and Justin M. Rao. 2016. Fair and Balanced? Quantifying Media Bias through Crowdsourced Content Analysis. *Public Opinion Quarterly*, 80(S1):250–271.

E. Cambria. 2016. Affective Computing and Sentiment Analysis. *IEEE Intelligent Systems*, 31(2):102–107.

Erik Cambria, Jie Fu, and Federica Bisio. 2015. AffectiveSpace 2: Enabling Affective Intuition for Concept-Level Sentiment Analysis. page 7.

Pew Research Center. 2018. Trends and Facts on Online News | State of the News Media.

Mark Coddington. 2015. Clarifying Journalism's Quantitative Turn: A typology for evaluating data journalism, computational journalism, and computer-assisted reporting. *Digital Journalism*, 3(3):331–348.

Claes H de Vreese. 2005. News framing: Theory and typology. page 12.

Devin M. Dotson, Susan K. Jacobson, Lynda Lee Kaid, and J. Stuart Carlton. 2012. Media Coverage of Climate Change in Chile: A Content Analysis of Conservative and Liberal Newspapers. *Environmental Communication*, 6(1):64–81.

Sven Engesser, Nicole Ernst, Frank Esser, and Florin Büchel. 2017. Populism and social media: How politicians spread a fragmented ideology. *Information, Communication & Society*, 20(8):1109–1126.

Sarah Eskens, Natali Helberger, and Judith Moeller. 2017. Challenged by news personalisation: Five perspectives on the right to receive information. *Journal of Media Law*, 9(2):259–284.

Lauren Feldman, Edward W. Maibach, Connie Roser-Renouf, and Anthony Leiserowitz. 2012. Climate on Cable: The Nature and Impact of Global Warming Coverage on Fox News, CNN, and MSNBC. *The International Journal of Press/Politics*, 17(1):3–31.

Flavio Figueiredo, Jussara M Almeida, Marcos Andre Gonc Alves, and Fabricio Benevenuto. 2014. On the Dynamics of Social Media Popularity: A YouTube Case Study. *ACM Transactions on Internet Technology*, 1(1):22.

Seth Flaxman, Sharad Goel, and Justin M. Rao. 2016. Filter Bubbles, Echo Chambers, and Online News Consumption. *Public Opinion Quarterly*, 80(S1):298–320.

Tomohiro Fukuhara, Hiroshi Nakagawa, and Toyoaki Nishida. 2007. Understanding Sentiment of People from News Articles: Temporal Sentiment Analysis of Social Events. page 2.

J. Gao, M. L. Jockers, J. Laudun, and T. Tangherlini. 2016. A multiscale theory for the dynamical evolution of sentiment in novels. In *2016 International Conference on Behavioral, Economic and Socio-Cultural Computing (BESC)*, pages 1–4.

Matthew Gentzkow and Jesse M. Shapiro. 2011. Ideological Segregation Online and Offline *. *The Quarterly Journal of Economics*, 126(4):1799–1839.

Jae Sik Ha and Donghee Shin. 2016. Framing the Arab Spring: Partisanship in the news stories of Korean Newspapers. *International Communication Gazette*, 78(6):536–556.

Michael Hameleers, Linda Bos, and Claes H. de Vreese. 2017. " *They* Did It": The Effects of Emotionalized Blame Attribution in Populist Communication. *Communication Research*, 44(6):870–900.

Natali Helberger. 2015. Merely Facilitating or Actively Stimulating Diverse Media Choices? Public Service Media at the Crossroad. page 17.

Matthew Jockers. 2015. » Revealing Sentiment and Plot Arcs with the Syuzhet Package Matthew L. Jockers.

Mesut Kaya, Guven Fidan, and Ismail H. Toroslu. 2012. Sentiment Analysis of Turkish Political News. In *2012 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology*, pages 174–180, Macau, China. IEEE.

Yaser Keneshloo, Shuguang Wang, Eui-Hong (Sam) Han, and Naren Ramakrishnan. 2016. Predicting the Popularity of News Articles. In *Proceedings of the 2016 SIAM International Conference on Data Mining*, pages 441–449. Society for Industrial and Applied Mathematics.

Bennett Kleinberg, Maximilian Mozes, and Isabelle van der Vegt. 2018. Identifying the sentiment styles of YouTube's vloggers. *In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, page 10.

Max Kuhn. 2008. Building Predictive Models in *R* Using the **caret** Package. *Journal of Statistical Software*, 28(5).

Bing Liu. 2010. Sentiment Analysis and Subjectivity. page 38.

Mika V. Mäntylä, Daniel Graziotin, and Miikka Kuutila. 2018. The evolution of sentiment analysis—A review of research topics, venues, and top cited papers. *Computer Science Review*, 27:16–32.

Maxwell E. McCombs. 2014. *Setting the Agenda: The Mass Media and Public Opinion*, 2. ed edition. Polity Press, Cambridge. OCLC: 935284613.

Rohit Mehta and Lynette DeAun Guzmán. 2018. Fake or Visual Trickery? Understanding the Quantitative Visual Rhetoric in the News. *Journal of Media Literacy Education*, page 19.

Eisa al Nashmi, Michael North, Terry Bloom, and Johanna Cleary. 2017. Promoting a global brand: a study of international news organisationsâĂŹ youtube channels. *The Journal of International Communication*, 23(2):165–185.

N. Newman, R. Fletcher, A. Kalogeropoulos, D. Levy, and R. K. Nielsen. 2018. Reuters Institute Digital News Report. http://www.digitalnewsreport.org/.

Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. page 94.

Limor Peer and Thomas B Ksiazek. 2011. Youtube and the challenge to journalism: new standards for news videos online. *Journalism Studies*, 12(1):45–63.

Leonard Richardson. 2019. Beautiful Soup.

Tyler Rinker. 2019a. A data package containing lexicons and dictionaries for text analysis: Trinker/lexicon.

Tyler Rinker. 2019b. Dictionary based sentiment analysis that considers valence shifters: Trinker/sentimentr.

Martijn Schoonvelde, Anna Brosius, Gijs Schumacher, and Bert N. Bakker. 2019a. Liberals lecture, conservatives communicate: Analyzing complexity and ideology in 381,609 political speeches. *PLOS ONE*, 14(2):e0208450.

Martijn Schoonvelde, Anna Brosius, Gijs Schumacher, and Bert N Bakker. 2019b. Liberals lecture, conservatives communicate: Analyzing complexity and ideology in 381,609 political speeches. *PloS one*, 14(2):e0208450.

Kim Christian Schrøder. 2015. News Media Old and New: Fluctuating audiences, news repertoires and locations of consumption. *Journalism Studies*, 16(1):60–78.

Saif Shahin, Pei Zheng, Heloisa Aruth Sturm, and Deepa Fadnis. 2016. Protesting the Paradigm: A Comparative Study of News Coverage of Protests in Brazil, China, and India. *The International Journal of Press/Politics*, 21(2):143–164.

Susan Currie Sivek. 2018. Both Facts and Feelings: Emotion and News Literacy. *Journal of Media Literacy Education*, page 16.

Stuart Soroka, Lori Young, and Meital Balmas. 2015. Bad News or Mad News? Sentiment Scoring of Negativity, Fear, and Anger in News Content. *The ANNALS of the American Academy of Political and Social Science*, 659(1):108–121.

Cass R. Sunstein. 2017. *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press, Princeton ; Oxford.

M. Iftekhar Tanveer, Samiha Samrose, Raiyan Abdul Baten, and M. Ehsan Hoque. 2018. Awe the Audience: How the Narrative Trajectories Affect Audience Perception in Public Speaking. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, pages 1–12, Montreal QC, Canada. ACM Press.

A. Tatar, P. Antoniadis, M. D. de Amorim, and S. Fdida. 2012. Ranking News Articles Based on Popularity Prediction. In *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 106–110, Istanbul. IEEE.

Robert L. Thorndike. 1953. Who belongs in the family? *Psychometrika*, 18(4):267–276.

Marc Trussler and Stuart Soroka. 2014. Consumer Demand for Cynical and Negative News Frames. *The International Journal of Press/Politics*, 19(3):360–379.

Peter Van Aelst, Jesper Strömbäck, Toril Aalberg, Frank Esser, Claes de Vreese, Jörg Matthes, David Hopmann, Susana Salgado, Nicolas Hubé, Agnieszka Stępińska, Stylianos Papathanassopoulos, Rosa Berganza, Guido Legnante, Carsten Reinemann, Tamir Sheafer, and James Stanyer. 2017. Political communication in a high-choice media environment: A challenge for democracy? *Annals of the International Communication Association*, 41(1):3–27.

Tyler Welch. 2018. Theology, heroism, justice, and fear: An analysis of ISIS propaganda magazines *Dabiq* and *Rumiyah*. *Dynamics of Asymmetric Conflict*, 11(3):186–198.

Lori Young and Stuart Soroka. 2012. Affective News: The Automated Coding of Sentiment in Political Texts. *Political Communication*, 29(2):205–231.