Identifying the Popularity and Persuasiveness of Right- and Left-Leaning Group Videos on Social Media

1st Lin Ai
Department of Computer Science
Columbia University
New York, U.S.A
lin.ai@cs.columbia.edu

2nd Anika Kathuria
Department of Computer Science
Columbia University
New York, U.S.A
ak4748@columbia.edu

3rd Subhadarshi Panda
Computer Science Department
Hunter College – CUNY
New York, U.S.A
sp2951@hunter.cuny.edu

4th Arushi Sahai
Department of Computer Science
Columbia University
New York, U.S.A
as5976@columbia.edu

5th Yuwen Yu
Computer Science Department
Hunter College – CUNY
New York, U.S.A
yyu4@gradcenter.cuny.edu

6th Sarah Ita Levitan
Computer Science Department
Hunter College – CUNY
New York, U.S.A
sarah.levitan@hunter.cuny.edu

6th Julia Hirschberg
Department of Computer Science
Columbia University
New York, U.S.A.
julia@cs.columbia.edu

***Abstract*—We have collected over 30,000 right- and left-leaning groups' videos from YouTube, Bitchute, 4Chan and Vimeo to identify aspects of their content and presentation which make these videos more popular and also potentially more persuasive. To date we have collected videos for and against Antifa and other anti-Fascist groups, Black Lives Matter, Proud Boys, Oath Keepers and QAnon and manually labelled subsets for style, stance toward the group, persuasiveness, techniques used and other features. We have also extracted video features including titles, descriptions, time of upload, captions and ASR transcripts, topic categories, and users' likes, dislikes, comments, and views. We are currently using these to automatically identify information such as the stance of the video (for or against a group), changes in popularity and in the sentiment of viewers toward the videos over time, correlating these changes with major events. We are also extracting text and audio features from videos and their comments to develop multimodal Machine Learning models for use in identifying different types of videos (e.g. pro- and anti- a group, extremely popular or unpopular) and eventually to use in identifying new radical groups and tracking their success. We will also be crowdsourcing surveys of subsets of these videos to understand how persons with different demographics and personality types perceive and are potentially influenced by different groups and different types of videos.**

***Index Terms*—radicalization, videos, far-right groups, far-left groups, popularity**

This work was funded by the Air Force Office of Scientific Research Award #FA9550-20-1-0400.

I. Introduction

Radicalization is the process of developing extremist ideologies and beliefs in others (Borum, 2011). In recent years, such efforts are increasingly seen on social media, where extremists spread their ideology and attempt to influence others to share extremist beliefs about race, ethnicity, gender, or religion through text-based posts, images, and videos. Previous work has proposed many theories of how and why radicalization develops, but little has been done to empirically test these theories on a large scale and to answer questions about specific features of radical messages beyond simple word lists that are significantly correlated with success in attracting followers. However, much of current radicalization is being attempted on sites like YouTube, Bitchute, Vimeo, and 4Chan through shared videos, or on groups' own websites, where many features beyond words can be used to attract viewers, including music, speech, and visual data. In this paper we describe current and ongoing research on identifying radicalizing groups' popularity through an analysis of their videos' metadata extracted from their online platforms. We have also begun using a variety of lexical features to identify a video's *stance* toward a group and will expand our feature set to include acoustic and visual features as our project proceeds.

II. Previous research

There has been considerable work by social scientists developing theories of how radicalization occurs and how movement leaders engage in cultural framing and cognitive interventions to build support [1]–[6]. Other researchers focus on how individuals' thoughts and behavior are influenced by the actual or imagined presence of others [7] and the process by which individual beliefs are transformed [8]–[10].

Mathematical game theoretic frameworks have also been proposed for how radicalization may occur in different scenarios [11]. There is also theoretical political science research on factors in social media that allow right-wing content creators to influence viewers by creating videos which are more impactful than text and comparing estimated viewing hours on far-right YouTube channels to estimated hours watching Fox News, CNN or MSNBC [12]. [13] also supports this view, examining the viewers of over 330,000 YouTube videos from 360 channels, mostly associated with far-right ideology, to identify viewers more likely to view right-wing videos.

However, there have been few other empirical studies that serve to support or explain social science theory. Most have used simple dictionary methods or sentiment analyses in social media to detect online radicalization [14]–[17]. Currently, more research is being done to identify features in other modalities, such as memes, using manual techniques applied to small datasets, in order to identify radicalization in visual images and videos [18]. Authors of [19] have used manually coded, qualitative judgments to identify visual aspects of social media to identify propaganda strategies of terrorist groups such as the Islamic State and Al Qaeda and its affiliates. Using a more linguistic perspective, Wilson's RedHen Lab [20] has been sharing data and methods and providing tools for collecting and tagging online resources to use in studies of radicalization. However, very little work has been done to develop more sophisticated computational methods for identifying the online multimodal elements which lead to radicalization.

III. Data collection

To create a large-scale corpus for our multimodal studies of online radicalization we have collected 31,673 videos from YouTube, BitChute, Vimeo, and Reddit for three right-leaning and 2 left-leaning groups. For right-leaning groups we have collected 5924 videos pro and anti QAnon; 1326 focusing on Proud Boys; and 589 on the Oath Keepers. For left-leaning videos we have collected 17,242 for Antifa (including 601 of these which are more general anti-Fascist videos) and 6592 for Black Lives Matter (BLM). We also collected number of views, likes, dislikes and comments for each group's videos. All of these videos are presented in English. QAnon, Oath Keepers, and Black Lives Matter videos were scraped from Youtube using the Youtube API's Search feature using the group's name as the search keyword. To collect videos from the BitChute platform, we used the name of each group as the keyword, and extracted the query results by scraping the HTML directly. The retrieved QAnon videos date from November 2017 to April 2021, the Proud Boys videos are from February 2018 to February 2021, the Oath Keepers videos are from December 2020 to February 2021, the Antifa videos are

Fig. 1. Q – Killing the Mockingbird.

Fig. 2. MESSAGE TO #ANTIFA FROM LOUISIANA LAW ENFORCEMENT.

from March 2017 to February 2021, and the BLM videos are from September 2019 to April 2021.

We selected these five groups because of their popularity in recent years and their active online presence. We chose both right-leaning and left-leaning groups in order to compare features of radicalization across the political spectrum. We note that the five groups are not equally extremist, and that there is considerable variation in extremism across groups and even within groups; we selected these as a basis for research and discovery of these differences. Below we describe the groups' main ideologies and a brief timeline of their popularity. Figures 1 and 2 display screenshots from initial clips of videos for two very different but far-right groups, QAnon[1] and Proud Boys[2], respectively.

QAnon is a far-right conspiracy theory alleging that a cabal of Satanic, cannibalistic pedophiles (which includes many famous Hollywood actors, Democratic politicians and government officials) is operating a global child sex trafficking ring while also conspiring against former President Trump during his term in office. It is often described as a cult. QAnon videos started picking up traction in the summer of 2018 after pictures of Q-supporters showing up to a Trump rally in August were released. Traffic started rapidly increasing after

[1]https://www.bitchute.com/video/uLBOYFLjce3R/
[2]https://www.bitchute.com/video/aIYLe0N1amBj/

the Covid-19 pandemic began, perhaps as people have had more time to watch videos, and there was a jump after January 6. Although YouTube banned QAnon videos in October 2020, their popularity has continued.

The Proud Boys is a far-right, male-only group that has engaged in political violence. The group was founded in 2016 in the midst of the presidential election and describes itself as "Western Chauvinists." It began gaining press after the 2017 Charlottesville car attack and there has been a steady increase in their videos' popularity until September 2020 when they were mentioned in a presidential debate; a sharp increase from then lasts till today with even more popularity after January 6, 2021.

The Oath Keepers, founded in 2009, claim that their organization includes tens of thousands of current and former law enforcement officers and military veterans and is one of the largest radical anti-government groups in the U.S. today. Like QAnon, it promotes a set of conspiracy theories claiming that the federal government is working to destroy American liberties. It began to gain traction in the summer of 2015, when the group made news by showing up to the Ferguson protests against the police shooting of Michael Brown in 2015 and increased significantly after the January 6 Capitol Riots.

The far-left groups are much less organized than far right groups but their videos have also attracted many views. Antifa is a left-wing anti-fascist movement in the US, and is largely decentralized with several independent subgroups. Its members tend to be anti-authoritarian, anti-capitalism, and anti-state and employ both violent and non-violent direct action to achieve their goals. Antifa began gaining popularity when the April 2017 far right and far left protests at UC Berkeley occurred, and Antifa was called out for sparking violence. Popularity increased in July 2019, when an individual who labeled himself an antifascist firebombed ICE and the present labeled Antifa as a terrorist organization. Group popularity grew until the May/June 2020 George Floyd protests began and there was much news focus on Antifa for looting and rioting.

Finally, Black Lives Matter (BLM) is another decentralized left-wing movement which began in protest of police brutality and other racially motivated violence against black people. BLM videos began increasing in July 2016 when many protests were held following the shootings of Alton in Baton Rouge, Philando Castilen in a St. Paul suburb, and Charles Kinsey in Miami. As with Antifa, BLM videos increased in posts until the end of May/beginning of June 2020. with more extreme growth since then, particularly after January 6 and in February 2021, possibly since this was Black History Month.

IV. Feature extraction

After collecting videos posted about the above five groups of interest, we automatically extracted a number of video features from the videos' metadata. The Youtube API was used to extract the features for Youtube videos - time of publication, numbers of likes, dislikes, views, related keywords, and comment text/likes/date of posting. Additionally, channel metadata was extracted(the name of the channel the video was posted on with additional information about the channel). We were also able to extract a number of text-based data from all of our videos including the video's title, its online description, captions, Automatic Speech Recognition (ASR) transcripts of the video, and viewer comments. We are currently extracting additional multimodal features including measurements of audio features (pitch, intensity, speaking rate, voice quality (jitter, shimmer, HNR) and background music and will eventually extract visual features as well, including facial expressions, gestures, and symbols (memes) used in the video.

Groups	QAnon	Antifa	Proud Boys	Oath Keepers	BLM
Total Videos	5,924	17,242	1,326	589	6,592
Comments	478,358	1,685,374	824,132	82,817	852,861
Avg Likes	255.88	575.82	7,051	790	23,217
Avg Dislikes	50.9	41.23	589	86	1,433
Avg Views	13,967	24,722	207,777	45,476	893,561

Fig. 3. Groups with Views, Likes, Dislikes and Comments.

Figure 3 summarizes the number of videos collected per group, along with metrics of video popularity, including the average number of views, comments, likes, and dislikes per group.

A. *Features used to identify popularity*

Our first goal was to identify automatic methods for identifying the popularity of different right- and left-leaning groups automatically from the video features extracted as described above. We decided to use number of views, number of comments and the sentiment expressed (positive/negative/neutral), and number of likes vs. dislikes. We also investigated how these changed over time and how they might be related to different major events. To do this we identified frequency of comments and likes over time, as shown in the graphs Figures 4-8. In almost all of these graphs, we see a peak around Sept 2020 when the U.S presidential debates began. The rise and fall of the peak associated with the number of likes are prominent, especially for Qanon and Antifa. However, for BLM the peak is on June 2020 and it continues to have around 200k likes per week in Jan 2021, which is the highest among all groups for that month. This suggests that the Capitol Riots not only inspired interest in far-right groups but also in their opposite, far left. The group with the lowest number of comments and of likes is the Oath Keepers. However, we also see that its trend is continuously increasing both in terms of the number of weekly comments and likes. Overall, the rise and fall of peaks for both QAnon and Antifa are the most significant, indicating the greater volatility in the popularity of their videos.

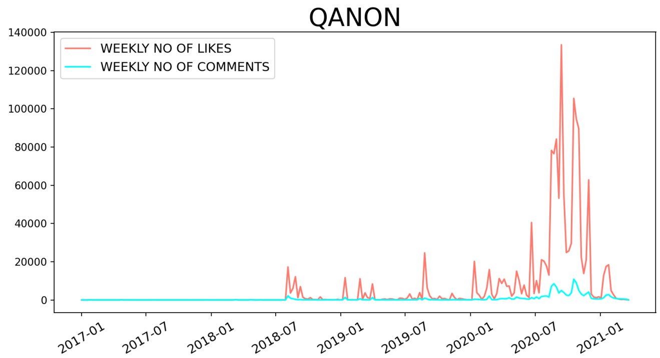

Fig. 4. QAnon Likes and Comments over Time.

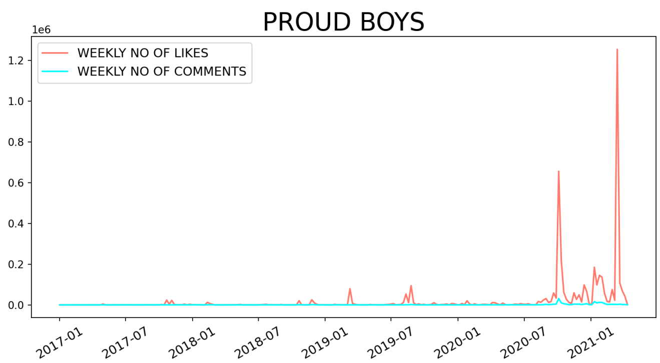

Fig. 5. Proud Boys Likes and Comments over Time.

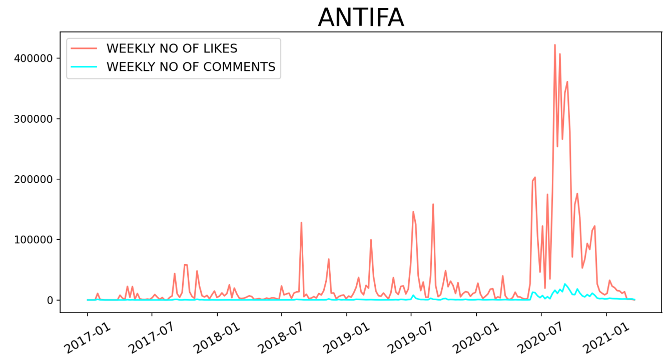

Fig. 7. Antifa Likes and Comments over Time.

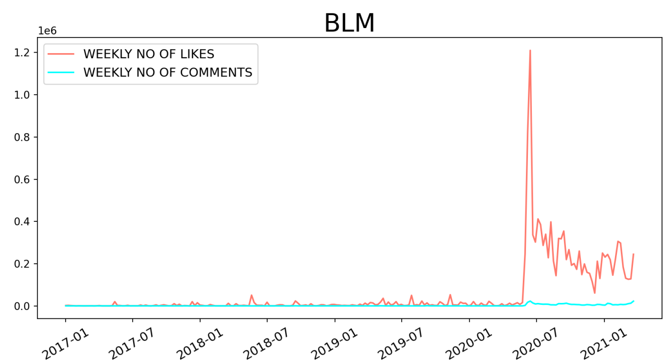

Fig. 8. Black Lives Matter Likes and Comments over Time.

Looking in more detail at what video comments can tell us for one of these groups, Antifa, we can see daily interaction in greater detail (Figure 9). In these graphs we see that Antifa videos and comments both began picking up traction April 2017, when, as noted earlier, far-left and far-right protests at UC Berkeley occurred and Antifa was called out for sparking violence. Comments also jumped in July 2019, was a self-proclaimed anti-fascist firebombed ICE and the president introduced his plan to label Antifa as a terrorist organization. There was steady growth until the end of May/beginning of June 2020, when the protests regarding George Floyd began and there was much focus on Antifa for looting and rioting. Since then, there has been a large increase in traffic. In 2021, though, while videos have continued increasing greatly, comments are not increasing in tandem.

V. Annotation

We have collected information available online from a large number of these videos, including the group the video discusses, whether it is a right- or left-leaning group, length of

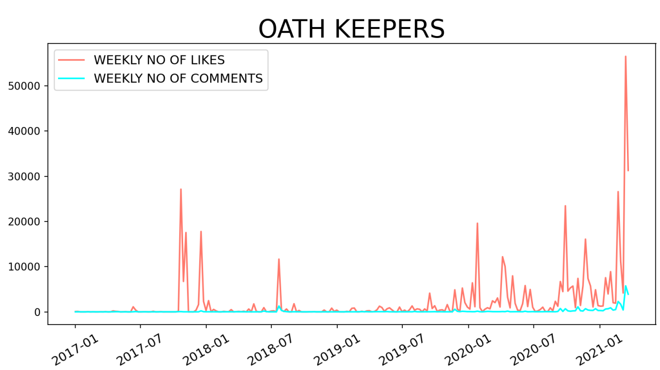

Fig. 6. Oath Keepers Likes and Comments over Time.

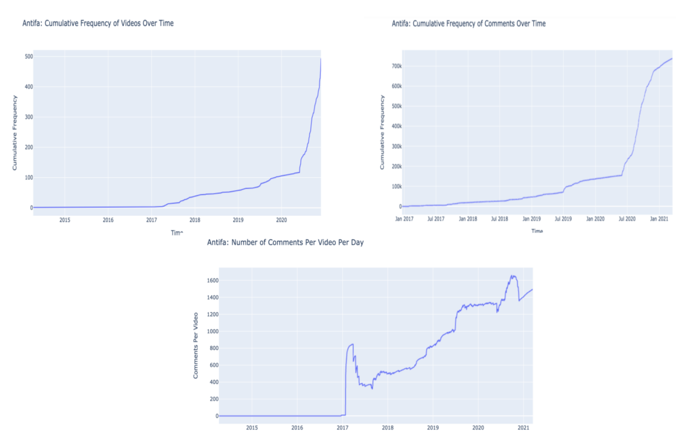

Fig. 9. Antifa Comments over Time.

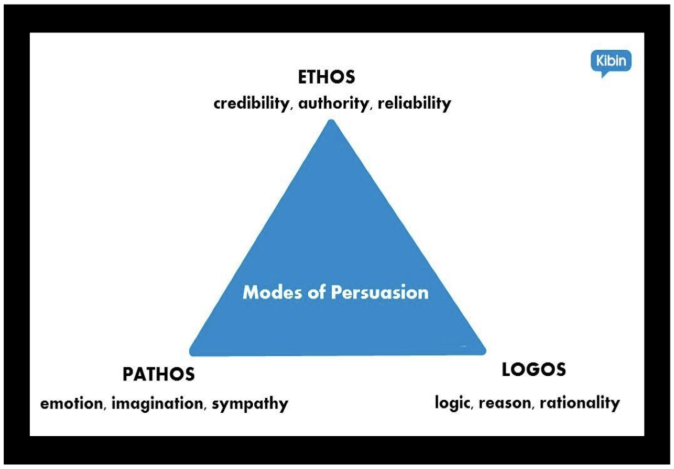

Fig. 10. The Rhetorical Triangle.

the video, the topic categories selected for the video from the platform list, tags used for YouTube videos to provide context for viewers, numbers of likes, dislikes, and comments, the channel which posted (and additional information about the channel's subscribers and topics. We also annotated a number of features ourselves for QAnon videos to identify the style of the video, whether the video includes music, voice- or text-over-visuals, other aspects of the video's style, such as vlogs (video blog posts), hosted videos, interviews, presentations, news broadcasts.

For a number of videos, we also annotated what we found to be the video's scores (1-3) on modes of presentation via the Rhetorical Triangle, including ethos (credibility, authority, reliability of content), logos (logic, reason, rationality), and pathos (emotion, imagination, sympathy). For a larger number of Antifa and QAnon we also annotated the videos *stance*, whether the videos was positive toward the group or negative, pro or con. We used these labels to develop models for automatic stance detection.

We have used these annotations to select subsets of the videos for further feature analysis and also to select videos that would be useful to put out for crowdsourcing to collect additional viewer information.

VI. Stance detection

One question we are using these annotation to answer is whether it is possible to identify whether a far-right or far-left group video is conveying positive, negative or neutral *stance* toward that group. To do this, we have selected a subset of our QAnon and Antifa videos which we annotated ourselves for stance. See Table I. We then used this subset to train Machine Learning classifiers on the rest of our data to develop a "weak-labeled" corpus for continuing research. We trained our ML model on a number of lexical features we extracted from the text associated with our QAnon and Antifa videos posted on YouTube, Bitchute and Vimeo.

To collect useful text-based features, we first pre-processed each video's title and description to remove URLs, punctuation, and stopwords (e.g. pronouns, prepositions) so that only words providing useful information for stance classification remained. We then word-tokenized and stemmed the remaining words using *Porter stemming* to reduce word inflections. We then performed Named Entity Recognition to identify potentially useful word types such as people, places, and organizations on the original, unprocessed data to retain capitalization but only performed this on video descriptions, since titles' capitalization would have made NER more difficult. We tagged words for part-of-speech to improve NER. We next identified *Bag of Words* features to identify the frequency of individual words (unigrams) in the data, normalizing using TF-IDF scores to prioritize words that are frequent but in relatively fewer videos to find the words most useful in identifying the stance of a video. Finally, we collected *Linguistic Inquiry and Word Count* (LIWC) information [21]. LIWC is a dictionary of words/stems labeled with 92 categories including psychological constructs, personal concerns, summary variables for analytical thinking, clout, authenticity and emotional tone.

Group	Pro/Anti	YouTube	Bitchute	Vimeo	Total
QAnon					
	Pro	0	100	N/A	100
	Anti	97	3	N/A	100
	Total	97	103	N/A	200
Antifa					
	Pro	97	5	34	136
	Anti	83	192	0	275
	Total	180	197	34	411

TABLE I
Stance Annotations on QAnon and Antifa Videos

Group	Pro/Anti	Prec	Rec	F1
QAnon				
	Pro	1.00	0.89	0.94
	Anti	0.91	1.00	0.95
Antifa				
	Pro	1.00	0.50	0.67
	Anti	0.80	1.00	0.89

TABLE II
Random Forest Precision/Recall/F1 on Test Set

We then used all these features to build ML models. We used Support Vector Machines and Random Forest methods for our models, since we wanted to obtain a good understanding of which features performed best on our data. We split our annotated subset for each group into 80 percent/20 percent training and test sets; for QAnon, we had 194 annotated videos and for Antifa, we had 358 (after pre-processing removed some videos from our original total). Results on our annotated test set are shown in Tables II and III.

Overall QAnon stance accuracies (generated through k-fold cross validation) were 91% (SVM), 94% (RF) and for Antifa 79% (SVM), 79% (RF). Based on these numbers, we selected the RF model for weak-labeling QAnon videos and the SVM for weak-labeling Antifa videos.

We also were able to identify a number of the most useful lexical features from our Random Forest and SVM models. These includes words like *White House, conspiracy* and *GOP*

Group	Pro/Anti	Prec	Rec	F1
QAnon				
	Pro	0.89	0.89	0.89
	Anti	0.90	0.90	0.90
Antifa				
	Pro	0.94	0.71	0.81
	Anti	0.87	0.98	0.92

TABLE III
SVM PRECISION/RECALL/F1 ON TEST SET

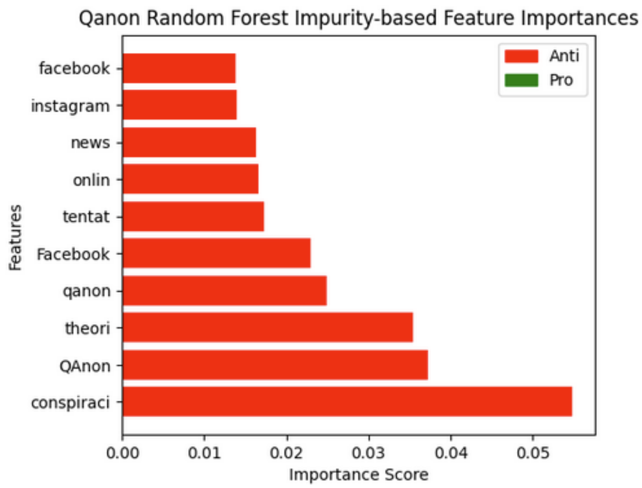

Fig. 11. Useful RF Stance Detection Features for QAnon

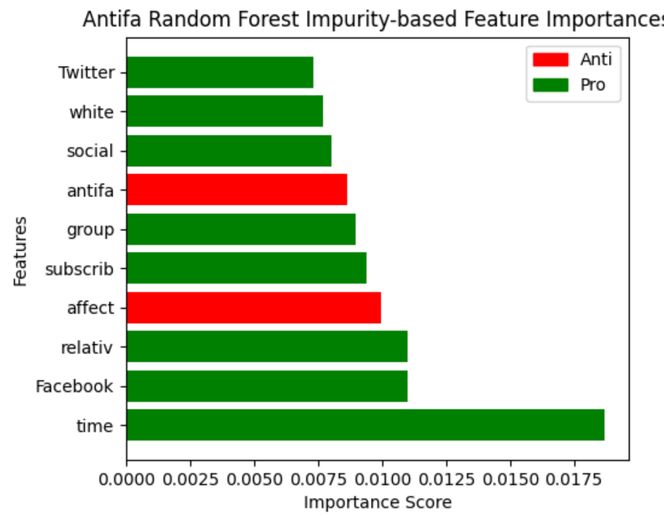

Fig. 12. Useful Stance Detection for Antifa Features

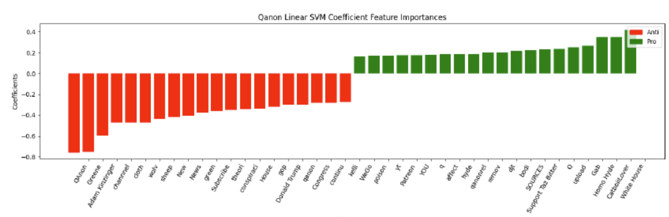

Fig. 13. Useful SVM Stance Detection Features for Pro- and Anti-QAnon Videos

for QAnon and words like *phoenix, CNN, Nazi,* and *Portland* for Antifa. A full list of features shown to be important for detecting which QAnon videos supported the group and which did not for the Random Forest models are show in Figure 11. Most of the features that were useful for Random Forests were negative features — which indicate that a video is against QAnon (in red). Features most useful for identifying stance using RF for our Antifa videos are show in Figure 12, where there are more features indicative of pro-Antifa videos (in green).

When we look at SVM features for QANon and Antifa videos there are more balanced pro- and anti-group features (shown in green for pro and red for anti). Figure 13 shows SVM Features useful for detecting pro- and anti-QAnon videos. Figure 14 shows SVM features useful for identifying pro- and anti-Antifa videos:

VII. Ongoing Research

Our next project will involve crowdsourcing tasks to ask viewers questions about what they found persuasive in our videos. For each task, we will provide short (˜3m) videos on one of our groups: half with be positive toward the group and half will be negative, based on our views and annotation. We will ask questions such as: How well do you think the video was produced? What were the most memorable aspects of the video — images, music, text, speech? How persuasive were the speakers in the video? Did the video make any valid points? What kind of reaction do you think other viewers might have to the video? We will use these ratings with other collected features to build ML models to identify more and less persuasive videos for current groups, as well as to detect potentially persuasive videos from new far-right or far-left leaning groups. We are also continuing work on improving the stance detection classifiers. We are planning to try BERT-style embeddings as well as Word2Vec embeddings.

VIII. Conclusion

The overall goal of our research is to improve our understanding of radicalization from online videos of left- and right-leaning groups, which aspects of online videos are most likely to appeal to viewers and incentivize them to appreciate and

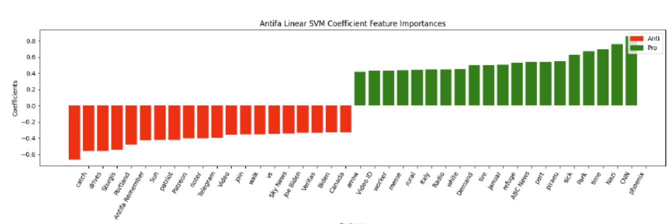

Fig. 14. Useful SVM Stance Detection Features for Pro- and Anti-Antifa Videos

perhaps even to join one of these groups. Toward this goal, we have collected over 30k videos from a variety of sources about 5 left- and right-leaning groups, annotated subsets of the videos, extracted features, and trained machine learning models for stance detection. We also analyzed the aggregated number of likes and comments per group over time, in order to understand how the popularity of the groups changes over time. To obtain more detailed information about how these videos may influence their viewers, we will crowdsource viewer opinions as well as collecting demographic information about them to determine what aspects of group videos are most persuasive and to which types of viewers. With this information we will continue to building Machine Learning models of potentially radicalizing videos in order to be able to identify similar videos computationally, using features extracted from the lexical, audio, and visual content of the videos. This could help to identify new, potentially radicalizing groups as well, through the videos they post. We also plan to make our very large collection of group videos publicly available to support research by others in radicalization.

ACKNOWLEDGMENTS

REFERENCES

[1] C. Geoge, "Regulating "Hate Spin": The Limits of Law in Managing Religious Incitement and Offense," International Journal of Communication 10 (2016), 2955–2972.

[2] C. George, Hate Spin: The Manufacture of Religious Offense and its Threat to Democracy. MIT Press: London, 2016.

[3] R. D. Benford and D. A. Snow, "Framing processes and social movements: An overview and assessment," Annual Review of Sociology, vol. 26, pp.611—639, 2000.

[4] W. A. Gamson, The Strategy of Social Protests, (2nd ed.). Belmont, CA: Wadsworth, 1990.

[5] W. A. Gamson, "Injustice frames," in D. A. Snow, D. Della Porta, B. Klandermans, and D. McAdam Eds., The Wiley-Blackwell Encyclopedia of Social and Political Movements (pp. 607–608), Oxford, UK: Blackwell, 2013. doi:10.1002/9780470674871.wbespm110

[6] R. Borum, "Radicalization into Violent Extremism I: A Review of Social Science Theories," Journal of Strategic Security, vol. 4, no. 4 (Winter 2011), pp. 7–36.

[7] C. McCauley and S. Moskalenko, "Mechanisms of Political Radicalization: Pathways Toward Terrorism," Terrorism and Political Violence, vol. 20, no. 3, pp.415–433, 2008, DOI: 10.1080/09546550802073367.

[8] L. R. Rambo, Understanding Religious Conversions, New Haven, CT: Yale University Press, 1993.

[9] Steinberg.

[10] T. Veldhuis and J. Staun, Islamist Radicalisation: A Root Cause Model, The Hague: Netherlands Institute of International Relations Clingendael, 2009.

[11] . M. B. Short, S. G.McCalla, M. R. D'Orsogna, "Modelling radicalization: how small violent fringe sects develop into large indoctrinated societies," Royal Society Open Science vol.4, no. 8, pp.170678, 2017.

[12] K. Munger and J.!Phillips, "Right-Wing YouTube: A Supply and Demand Perspective," The International Journal of Press/Politics, 2020, https://doi.org/10.1177/1940161220964767.

[13] M. Ribeiro, R. Ottoni, R. West, V. A. F. Almeida and W. Meira, "Auditing Radicalization Pathways on YouTube," Proceedings of the 2020 Converence on Fairness, Accountability, and Transparency, January 2020, pp.131-141.

[14] S Agarwal, "Applying Social Media Intelligence for Predicting and Identifying On-line Radicalization and Civil Unrest Oriented Threats," Literature Review (PhD comprehensive exam), arXiv: 1511.06858v1, 11/21/2015.

[15] S. Agarwal and A. Sureka, "Topic-Specific YouTube Crawling to Detect Online Radicalization," International Workshop on Databases in Networked Information Systems, pp. 133–151, 2015.

[16] T. Grover and G. Mark, "Detecting Potential Warning Behaviors of Ideological Radicalization in an Alt-Right Subreddit," Thirteenth International AAAI Conference on Web and Social Media (ICWSM 2019), pp. 193-204.

[17] I. van der Vegt, M. Mozes, P. Gill, and B. Kleinberg, "Online influence, offline violence: Linguistic responses to the 'Unite the Right' rally," Journal of Computational Social Science, vol.4, pp.333–354, 2021.

[18] Am. Acker and J. Donovan, "Data Craft: A Theory/Methods Package for Critical Internet Studies." Information, Communication and Society vol.22, no.11, pp.1590-1609, 2019.

[19] M. D. Robinson and C. E. Dauber, "Grading the Quality of ISIS Videos: A Metric for Assessing the Technical Sophistication of Digital Video Propaganda," Studies in Conflict and Terrorism, vol. 42, nos.1-2, pp.70-87, DOI: 10.1080/1057610X.2018.1513693.

[20] F. F. Steen, A. Hougaard, J. Joo, I. Olza, C. PagánCánovas, A. Pleshakova, S. Ray, P. Uhrig, J. Valenzuela, J. Woźny and M. Turner, "Toward an Infrastructure for Data-driven Multimodal Communication," Linguistics Vanguard, vol.4, no.1, pp.1-9, March 2018.

[21] J. Pennebaker, R. L.Boyd, K. Jordan, and K. Blackburn, The Development and Psychometric Properties of LIWC2015, Austin,TX: University of Texas at Austin, 2015.