Amsterdam University Press

COMPUTATIONAL COMMUNICATION RESEARCH 5.1 (2023) 1–33
HTTPS://DOI.ORG/10.5117/CCR2023.1.11.GITO

The speech we miss: How keyword-based data collection obscures youth participation in online political discourse

Adina Gitomer
Network Science Institute, Northeastern University, Boston, USA

Sarah Shugars
School of Communication and Information, Rutgers University, New Brunswick, USA

Ryan J. Gallagher
Network Science Institute, Northeastern University, Boston, USA

Stefan McCabe
Institute for Data, Democracy & Politics, The George Washington University, Washington, D.C., USA

Brooke Foucault Welles
College of Arts, Media, and Design, Northeastern University, Boston, USA

Abstract
In this work, we leverage a panel of over 1.6 million Twitter users matched with public voter records to assess how a standard keyword-based approach to social media data collection performs in the context of participatory politics, and we critically examine the speech this method leaves behind. We find that keyword classifiers undercount young people's participation in online political discourse, and that valuable political expression is lost in the process. We argue that a mainstream keyword approach to collecting social media data is not well-suited to the participatory politics associated with young people and may reinforce a false perception of youth political apathy as a result.

Keywords: keyword classifiers, participatory politics, online discourse

Introduction

Young people's political participation in the United States is in many ways devalued. Citing low scores on traditional participation measures, youth are often regarded as politically apathetic, uninformed, and precipitating a broader crisis in democracy (Flanagan & Levine, 2010). While this degrading view of youth has persisted in scholarly literature and the popular imagination alike, it is increasingly contested. Youth face structural barriers to institutional forms of participation, as US political institutions were neither created nor advanced with them in mind (Zhu et al., 2019). Meanwhile, having grown up with new media technologies and the ability to give direct input and feedback in real time, youth have developed new sensibilities for what participation looks like (Jenkins et al., 2015). These two phenomena have led youth to engage in a "participatory politics" that aims to shape political outcomes by emphasizing peer interaction and discursive power over formal elites and institutions (Cohen & Kahne, 2012; Jenkins et al., 2016). Such foci render traditional political participation metrics (e.g., voter turnout) inadequate, and suggest that relying on these metrics misrepresents the extent of youth political participation. To fully analyze and appreciate youth political participation, then, requires investigating participatory politics directly.

Political discourse on social media is a primary site of participatory politics, and a natural context for studying its form and function. To begin, social media is the infrastructure that has brought participatory politics to life: It enables the amplification of voices outside of mainstream media (Jenkins et al., 2016), and creates the space to connect, share, and deliberate through peer-to-peer networks (Cohen & Kahne, 2012). Beyond that, as mentioned, discourse is a key lever in the participatory politics framework. It is central, among other reasons, because it doesn't rely on institutional access, and can prompt the kind of shifts in public attitudes that have historically accomplished more for young people than electoral victories (Jenkins et al., 2016). While more traditional definitions of political participation may not include discursive engagement, favoring instead behavioral acts like voting and canvassing, the participatory politics literature explicitly positions it as an active and effective form of participation (Van Deth, 2016).

From a data perspective, turning to political discourse on social media offers the potential to analyze participatory politics at scale and through its direct artifacts. To realize that potential, however, requires accurately identifying those artifacts in the first place; failure on this front could mean continuing the harmful legacy of using unfit measures to judge young people's political participation. In this work, we leverage a panel that matches over 1.6 million Twitter users with public voter records to assess how a standard keyword-based approach to social media data collection performs in the context of participatory politics. Focusing on the 2020 US presidential election, we find that keyword classifiers, whether expert-curated or inductively-defined, obscure and undercount young people's (20-29) participation in the election discourse.

Moreover, through a critical discourse analysis, we discover that the speech overlooked by a keyword method contains rich discussion of the election. Humor and emotional appeals feature prominently therein, where they signal political awareness, mediate the construction of in- and out-groups, and renegotiate political power by breaking with mainstream conventions of political debate. To miss this speech, then, is to miss important contributions to the online election discourse and powerful examples of youth political participation more broadly. We ultimately argue that a mainstream keyword approach to collecting social media data is not well-suited to studying the participatory politics associated with young people and is poised to reinforce long-standing political power imbalances as a result.

Related Work

Participatory Politics

Young people's political participation is often discussed through a narrative of decline; as Jenkins et al. (2016) explain, young people are "seen as emblematic of the crisis in democracy—represented as apathetic about institutional politics, ill-informed about current affairs, and unwilling to register and vote" (p. 7). This perception is supported by a range of findings and arguments tied to traditional ways of measuring political participation. To name a few: youth make up the smallest proportion of voters (Khosla, 2022; O'Neill, 2022), are increasingly detached from political institutions (Putnam, 2000; Sloam & Henn, 2019), tend to believe their political involvement is inconsequential (Harvard Youth Poll: Top Trends and Takeaways, 2022), exhibit low levels of news consumption and political knowledge (Wattenberg, 2015), and express low confidence in governmental institutions (Pew Research Center, 2022).

While the findings listed above may be true, the conclusion that youth are politically apathetic is not. As Bennett et al. (2011) argue, a lack of participation according to these measures points to the "fragmentation of an old civic order"; meanwhile, young people are at the forefront of "emerging civic styles." Young people's turn toward new forms of participation is hardly surprising, since they face a number of barriers in the formal political sphere. For example, there are minimum age restrictions on voting, and even stricter ones on running for public office (Longley, 2022; Nwanevu, 2014); further, youth are less likely to be stable homeowners, which has a negative effect on electoral participation (Mejia et al., 2018). Compounding all of this, the issues they find important tend to be ignored by elected officials (King & Weisman, 2022).

The emergent political style associated with youth has been named "participatory politics," which Cohen and Kahne (2012) define as "interactive, peer-based acts through which individuals and groups seek to exert voice and influence on issues of public concern." Further, "these acts are not guided by deference to elites or formal institutions." As such, participatory politics "facilitates a renegotiation of political power," and lends more agency and independence to non-elite political actors.

Participatory politics has been framed as an outgrowth of a more general "participatory culture," wherein everyday individuals are encouraged to take part in media making practices that they find meaningful, often inspired by cultural interests (Jenkins, 2006). Participatory culture emphasizes creating and sharing content, fueled by a sense of empowerment over that process and the ability to forge social ties through it (Jenkins et al., 2017). The recent proliferation of new media technologies has led to a "more" participatory culture, in which the majority of youth today have been socialized (Prensky, 2001). These youth have increasingly employed the "skills norms, and networks" (Cohen & Kahne, 2012) of participatory culture for political ends, thus giving way to participatory politics. This lineage positions participatory politics as a playful, confident form of political engagement that centers community bonds, "social and cultural mechanisms," and "expressive and discursive power" (Jenkins et al., 2016, pp. 2, 4). Moreover, it primarily plays out on social media, where all users are de facto media makers, and creative personal expression reigns supreme.

Given all of the above, participatory politics can be seen as a turn away from the hierarchical model that characterizes American political institutions, and a step toward a more inclusive and empowering system. Indeed, despite digital divides and inequalities, participatory politics has been found to be equitably distributed across demographic groups (Cohen & Kahne, 2012). Recent studies have upheld this finding, showing that online political participation attracts youth from more diverse backgrounds and is less stratified by social economic status than traditional forms of participation, like voting (Lane et al., 2023; Vromen et al., 2016).

Though it remains a departure from some longstanding notions of political participation, recent work has framed youth participatory politics as essential to social and political processes (Kligler-Vilenchik & Literat, 2020), and has increasingly taken its legitimacy as a starting point (Lane, 2020). This has allowed participatory politics research to move beyond proving its significance and, primarily by examining youth political expression online, focus on exploring its contours and complexities. For instance, scholars have found that while political memes and pop culture references are central, they are not uniformly loved or viewed as beneficial (Literat & Kligler-Vilenchik, 2021; Penney, 2020). At the same time, political agency and belonging have been definitively located within these creative strategies (Kligler-Vilenchik & Literat, 2018). Other work has detailed the way digital affordances figure in, including how connection to like-minded others tends to support youth political expression online (Literat & Kligler-Vilenchik, 2019), while design choices like identifiability tend to limit it (Lane, 2020). As a whole, these studies have helped to paint a more nuanced and hybridized picture of participatory politics (Penney, 2019), while maintaining its worth and political pertinence.

Role of New Media and Data Collection

As is evident from the historical and theoretical underpinnings of participatory politics cited above, as well as the recent body of work on the topic, young people's online political discourse is a promising site for interrogating the nature and application of participatory politics. Going directly to the source, it appreciates the prevalence of exerting voice, interacting with peers, creating content, and circulating ideas. This approach also allows for large-scale analyses, and reference to participatory politics' direct artifacts in the form of the social media posts themselves.

When building a corpus of social media discourse around a certain topic, political or otherwise, it is extremely common for researchers to determine discursive relevance through the use of a keyword classifier (Chen et al., 2022; Driscoll & Walker, 2014; King et al., 2017; Pew Research Center, 2019). This involves generating a list of words that pertain to the topic of interest and pulling all the posts that include at least one of those words. The process is entirely automated, as the universe of posts on social media is far too vast for manual sorting. It is also relatively straightforward and intuitive, and more expansive than, say, collecting data solely based on hashtags, as it attempts to capture a thematic conversation in full.

Between the potential of studying observational social media data to understand participatory politics, and the legacy of using unfit metrics to evaluate youth political participation, this paper asks: **How does a keyword-based classifier perform in the context of young people's online 2020**

presidential election discourse? What kind of speech is missed through this approach? How does the speech we miss matter for purposes of measurement and our ability to quantify participation in online political discourses? How does it matter in terms of our normative understandings of participatory politics and the boundaries of political participation?

Taken together, the answers to these questions illuminate the ways in which a standard methodological choice can compromise our measurements of and appreciation for political participation – particularly among youth. In pursuit of these answers, we focus on the 2020 presidential election because it was a major event in US politics, receiving extensive attention across demographic groups, including youth (Penney, 2019).

Data

Twitter Panel

To address our questions, we drew from a panel developed by Grinberg et al. (2019) that matches 1,643,182 Twitter users with public US voter records. The panel was created by sampling users from Twitter's 10% Decahose between 2014 and 2017 and matching those users on name and location with public voter records compiled by the data vendor TargetSmart in 2017. More details on the exact matching process can be found in Grinberg et al., 2019; Shugars et al., 2021; and Hughes et al., 2021. Because the voter records are from 2017 (and include a share of preregistered 17-year-olds), the youngest users in the panel were 20 years old during the 2020 US presidential election.

Matching Twitter users to voter records ensures that we are studying real individuals, as opposed to organizations or bots (Gorwa & Guilbeault, 2020; Varol et al., 2017). More importantly, we can associate these individuals with demographic attributes detailed in their voter files, including age. That said, the panel contains biases based on which Twitter users can be successfully matched by name and location. Prior work shows that our panel slightly over-represents white users and women, and undercounts Asians and Hispanics (Hughes et al., 2021). In addition, non-registered citizens and disenfranchised people make important contributions to election discourse, but they are excluded from the panel by design (Kennedy & Deane, 2017; White & Nguyen, 2022). Others still are simply unlisted in voter files and similar data (Jackman & Spahn, 2021). Despite these limitations, our panel is a powerful dataset for studying participatory politics via social media discourse.

We chose to focus on Twitter because it has an undeniable effect on the

American political climate (Munger, 2017), plays an instrumental role in organizing political movements globally (Lotan et al., 2011; Tufekci, 2017), and may be shifting the balance of political visibility (Freelon et al., 2018; Jackson et al., 2020). Moreover, Twitter sees itself as playing “a critical role” in “empowering democratic conversation [and] driving civic participation,” and took concrete steps to modify and augment its platform in anticipation of the 2020 presidential election (Gadde & Beykpour, 2020). Twitter’s actions were especially reasonable, as the election is a news event, and Twitter is a main news site for many Americans (Mitchell, et al., 2021). While Twitter is not the primary platform among young people, 18-29 year-olds are more likely than any other cohort to have an account (Pew Research Center, 2021).

Election Tweets

We collected all panelist tweets between November 1st and 15th of 2020, which we set as our election window. This resulted in a total corpus of 18,634,163 tweets generated by 418, 976 unique users across all age groups. In general, users aged 30-49 were the most prolific, having produced 38% of the tweets in the corpus; users aged 20-29 produced 18% of the tweets, those aged 50-64 produced 29% of tweets, and those aged 65-89 produced 14% of tweets.

Since we were dealing with a large number of tweets, we took a sample of them for hand coding. Because we wanted our sample to capture both election and non-election tweets, we first employed a keyword-based classifier to label every tweet in our initial corpus as either election related or not. As part of our analysis, we inspected the accuracy of this classifier and determined whether a better keyword-based classifier could be constructed. However, for the purpose of sampling our data, this initial classifier helped ensure that our final corpus of hand-coded tweets had some representation of both election and non-election content. Specifically, tweets were marked as election related by the classifier if they matched one of our 118 expert-curated election keywords representing a range of election discourse (see Appendix). Tweets were also marked as election related if they contained a relevant URL (Gallagher et al., 2021; Shugars et al., 2021), where a URL is considered relevant if it co-occured with keywords at least 100 times and at least 20% of its use was with at least one keyword; these cutoffs were chosen because they improve the sensitivity of identifying topical URLs without compromising precision (Shugars et al., 2021). The resulting election corpus contained 6,428,620 tweets – 34% of the total set; the remaining 66% of tweets were labeled as non-election related.

Next, for each individual age in a cohort, we sampled the top election and non-election tweets according to retweet count. Specifically, we took the top 0.4% of tweets, or the top 100 tweets if 0.4% yielded less than that. We focused on the most popular (i.e., most retweeted) tweets – a common approach in the case of Twitter data (Jackson et al., 2020) – since they accrued more visibility and likely held more discussive influence than the average tweet. While they may not be representative of election tweets writ large, which would have required a random sampling strategy, they typify those most prominent within the election discourse. The sampling thresholds were chosen to ensure the sample sizes were substantially large and representative, as well as feasible to read given our time and resources. We incorporated the absolute value (100) in addition to the percentage (0.4%) because the variance in the total number of election tweets per age occasionally rendered the percent threshold insufficient; however, when possible, we preferred the percentage for its ability to conserve the distribution of tweets by age found in the full dataset. After aggregating individual ages back into four cohort brackets (20-29, 30-49, 50-64, and 65-89), we were left with a total of 8 unique samples, the sizes of which are reported in Table 1; each sample covers over half of the retweets for the corresponding age group in the initial corpora. We used the composite sample of 16,245 tweets as our primary corpus for analysis.

Cohort	Election Sample (N)	Non-Election Sample (N)	False Negative Rate	False Positive Rate	Precision	Recall	Accuracy	F1 Score
20-29	883	1856	36.6%	3.44%	0.945	0.634	0.806	0.759
30-49	2084	2941	39.3%	5.43%	0.954	0.607	0.725	0.742
50-64	1580	1768	25.1%	6.22%	0.947	0.749	0.825	0.836
65-89	2565	2568	16.1%	8.30%	0.925	0.839	0.874	0.880

Table 1: Tweet Sample Sizes and Classifier Accuracy Per Cohort

Methods

We hand coded every tweet in the sample (N = 16,245) for whether it is election related or not, giving us gold-standard labels against which to measure classifier accuracy. Two types of election tweets were flagged: 1) tweets that are plainly related to the presidential election through salient text or visuals, and 2) tweets that are vaguely related to the presidential election, in that they seem to strongly assume an audience primed by election events. For more details on the exact coding procedure, refer to the codebook in the Appendix. The hand coding was carried out by three coders, who first coded a reliability sample, reconvened to discuss and update the codebook, and then recoded the sample. The three coders achieved a percent agreement of 0.965, a Brennan-Prediger coefficient of 0.931, and a Krippendorff's Alpha coefficient of 0.921.

Classifier Accuracy

We began addressing our first research question by assessing the accuracy of keyword-based classifiers. In addition to examining the performance of our expert-curated keyword classifier described in the previous section, we used our labeled corpus to inductively construct and assess the performance of numerous keyword classifiers. Specifically, we split our data into an 80% training set and 20% test set, using age information and hand-coded labels to balance by both age and election content. We then used information about word frequency in the training set to construct 10,000 potential keyword classifiers (i.e., keyword lists), which we assessed with the test set.

We constructed these classifiers using a 2-parameter model: For each unique word in the training corpus, it was included in the keyword classifier if it (1) occurred within the training corpus more frequently than some minimum occurrence threshold α, and if (2) the ratio of occurrence in election tweets vs. non-election tweets was greater than some threshold β. In other words, the first parameter helps control for rare words by requiring a keyword to appear at least α times, while the second parameter helps control for common words by requiring a keyword to be β times more frequent in election-related tweets.

We used the hand labels to identify a tweet as election related or not. We searched 100 threshold values, from 0-99, for each parameter, resulting in a total of 10,000 (100×100) potential classifiers. For both parameters, the lower threshold of 0 represents a naïve baseline in which the parameter does not restrict any words from being used in the classifier. The upper threshold of 99 represents an expected reasonable maximum, with α = 99 meaning that a term must occur at least 99 times to be considered "not rare" and β = 99 meaning that a term must occur 99 times more in the election corpus to be considered "not common." Together, these parameters map the space of potential keyword classifier performance.

False Negative Analysis

We were particularly interested here in false negatives – tweets which are election related, but which were computationally labeled as non-election related – since they represent what a keyword classifier *overlooks*. This analysis was used to continue addressing our first research question of how keyword classifiers perform in the face of participatory politics, and to begin addressing our second research question of what we miss under such an approach. It was conducted over the set of false negatives generated by our expert-curated keyword classifier.

Since keyword methods rely exclusively on text, we looked at the distribution of media attachments in the set of false negatives to see if media use was driving the high false negative rates. Additionally, we inductively enhanced our expert-curated keyword list by examining the 30 most frequently used terms in each cohort's set of false negatives, and manually adding those that would make for good election keywords. The first author made the initial selection, informally pulling out all words with a clear tie to the election and ignoring more general and/or ambiguous terms (e.g., "tonight", "right"); following this, the rest of the authors reviewed and endorsed the selection. This process generated a list of 10 new keywords: "president," "votes," "voters," "ballots," "democracy," "republicans," "donald," "four years," "4 years," "four seasons." Next, we ensured that this keyword list produced fewer false positives than the original list for every individual cohort (its maximum rate was 2.2% for 50-64 year-olds, which is more than a point below the minimum rate in the original classifier, as per Table 1). Although this confirmed that the keywords in the inductively-expanded list did not heighten the issue of false positives, the blatantly election-related terms therein (e.g., "president") were not included in the original list over concerns that they would do just that. Finally, we measured the proportion of "keywordifiable" false negatives per cohort, or those that would be correctly identified as election related, had these inductively-determined keywords been included in the initial list. That is, we examined the share of election tweets that contain election-related keywords (i.e., are "keywordifiable"), but that were overlooked by the original classifier because those particular keywords were not included in the original list.

Figure 1 summarizes the various procedures we used to classify election tweets in this work, and how they are related to each other.

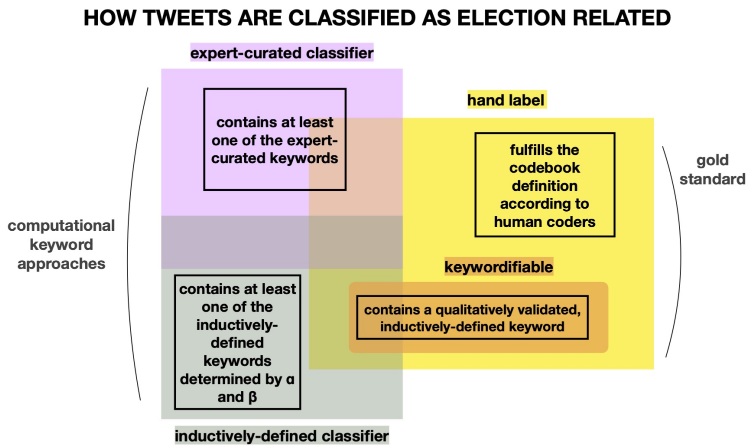

Figure 1: Schematic of Election Tweet Classification Procedures

Critical Discourse Analysis

Finally, to continue addressing what is lost through a keyword approach, and to determine how that loss matters both normatively and in terms of measurement, we conducted a critical discourse analysis (Fairclough, 2013) on the set of young user's election tweets that were mislabeled by our expert-curated classifier. We critically analyzed the tweet content, while centering users' political positionalities and relationships with power. Our goal was to unpack the contents of the overlooked election speech, elaborate on the challenges it poses to a keyword approach, and spell out how it is meaningful from a participatory politics lens.

Results

Classifier Accuracy

Our initial classifier based on an expert-curated list of 118 keywords achieved a performance (F1 score) of 80.9%, which is highly consistent with the best performance achieved (80.3%) by our inductive search over potential keyword lists. Across all models, keyword classifiers tended to have low false positive rates but high false negative rates. In other words, tweets labeled as election related by a keyword classifier most likely are on topic, but these classifiers potentially missed a large number of election-related tweets, erroneously identifying them as not related to the election. Since our sample corpus included the most popular (i.e., retweeted) content, this suggests that these classifiers miss significant portions of the dialogue.

Table 1 shows how this dynamic breaks down across age cohorts, reporting key measures of performance for the expert-curated classifier. Here we see that the false positive rates were considerably low across age groups: When the classifier labeled a tweet as election related, it was usually correct. However, these rates were higher for older users, meaning the classifier was more likely to exaggerate their contributions to the election discourse. Meanwhile, the false negative rates (see Table 1) were high across the board, suggesting that the keyword classifier often missed tweets that truly are election related. False negatives were more common for the youngest cohorts (36.6% & 39.3%) than for the oldest cohorts (25.1% & 16.1%). Thus, when confronted with an election tweet, there was a large chance the classifier would declare it unrelated to the election – especially if it was posted by a younger user.

We saw similar results in our inductive search of potential keyword lists, with classifiers consistently under-representing the election content produced by younger users. Using this approach, our best performing classifier required keywords to occur within the training corpus at least 83 times and to occur 10 times more frequently in election tweets than non-election tweets. See Figure 1 in the Appendix for classifier performance across all parameter values. Similar to our expert-curated classifier, this model achieved an F1 score of 80.3% for all users, and generally performed better for older users than younger users. Namely, this model achieved an F1 score of 83.3% for those 65-89, an F1 score of 80.8% for those 50-64, an F1 score of 79.5% for those 30-49, and finally an F1 score of 73.3% for those 20-29.

Again, high false negative rates seem to be driving this disparity. Figure 2 shows how classifiers' false negative rates per cohort varied based on the ratio of occurrence between election and non-election tweets. To better visualize this variation in performance across our β parameter, Figure 2 fixes the minimum occurrence parameter to $\alpha = 83$, the value which achieved the best performance across all α parameters. The resulting classifiers across varied β parameters consistently produced more false negatives for the youngest users (20-29) and fewer false negatives for the oldest users (65-89). Particularly bad classifiers also had high false negative rates for those 30-49, but even the best possible keyword-based classifiers had poor performance for 20-29 year-olds, while maintaining reasonably good performance for other age cohorts. Across the 10,000 keyword-based classifiers we tried, the model which gave the best performance for 20-29 year-olds still falsely labeled nearly a quarter (23.1%) of their election tweets as non-election related. This performance is notably worse than that of other age groups: The best classifier for 30-49 year-olds mislabeled 1/5 of their election tweets (20.7%) as false negatives, while the best classifiers for 50-64 year-olds and 65-89 year-olds both only had a false negative rate of 9.7%.

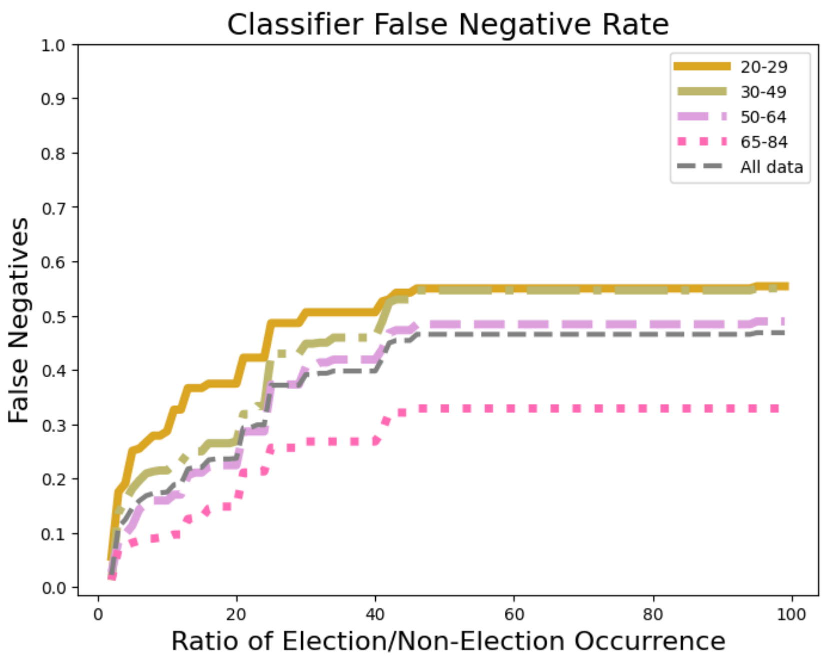

Figure 2: The false negative rate per cohort (and overall) as the ratio of occurrence in election tweets to non-election tweets threshold parameter (β) varies from 0 to 99. The minimum occurrence parameter (α) is fixed to 83, the value which had the best performance. The youngest cohort (20-29) maintains the highest rate of false negatives across all β's, while the oldest cohort (65-89) maintains the lowest rate.

Together, these results suggest that keyword-based classifiers systematically overcount older users' tweets and undercount younger users' tweets. The low level of false positives across models and age cohorts further suggests that this effect is primarily driven by high false negative rates: For younger users in particular, keyword-based models miss a substantial amount of election-related speech.

False Negative Analysis

Figure 3 shows the media breakdown of each cohort's set of false negatives. While users of all ages included media in their election tweets, 20-29 year-olds were the most media-happy: Over half of their mislabeled tweets contained a media attachment. This is significant, as most keyword classifiers – ours included – are built to work with text, and do not deal with media at all. Although we may be able to pull text out of media objects, this is not the norm for keyword approaches, it greatly increases the complexity of the classification task, and the election cues still may not exist in the form of text (for example: A picture or video of a candidate).

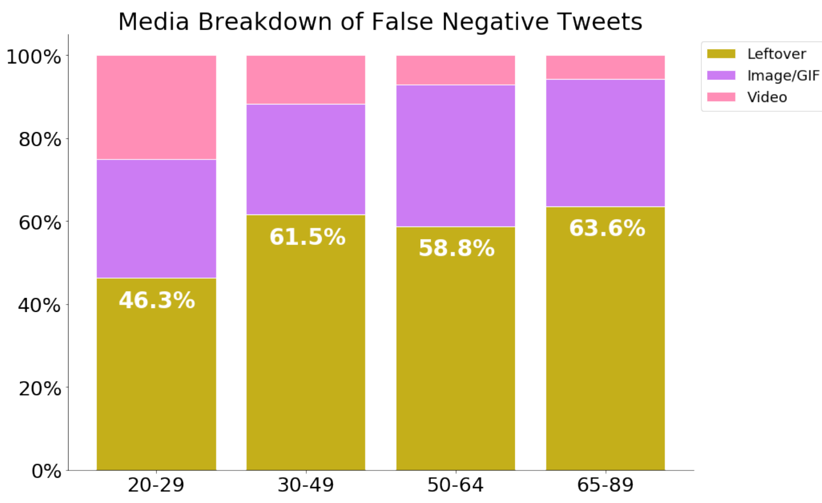

Figure 3: The proportion of images/gifs and videos for each cohort's set of misclassified election tweets. The "leftover" segments – labeled with their exact percentages - represent the proportion of misclassified election tweets that do not contain any media attachments.

Next, Figure 4 visualizes the proportion of "keywordifiable" false negatives per cohort, or the proportion of false negatives that would have been correctly labeled in the case of an inductively expanded keyword list. As previously described in the methods section, this classifier includes the 118 expertly-curated keywords listed in the Appendix, along with a manually-curated list of terms which occurred frequently in false negative election tweets. The youngest cohort had the smallest proportion of keywordifiable tweets, at 11.2%. Chi-squared tests confirmed that 20-29 year-olds had statistically significantly ($p < .002$) fewer keywordifiable tweets than would be expected when compared with every other cohort[1]. Building on the results in the previous subsection, this indicates that younger users employed more non-standard and varied language to talk about the election on Twitter, and an improved keyword list would not have greatly improved our ability to detect their election speech.

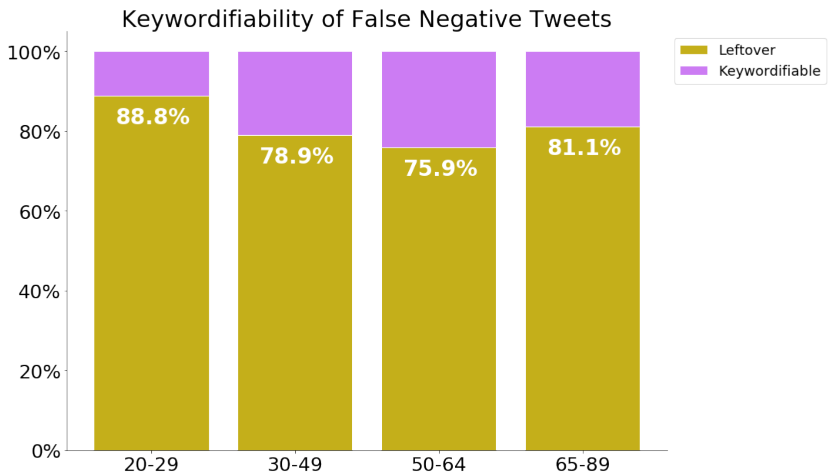

Figure 4: The proportion of each cohort's misclassified election tweets that are "keywordifiable," or that contain a keyword match with our manually updated keyword list. The "leftover" segments – labeled with their exact percentages – represent the proportion of misclassified election tweets that do not contain a keyword match with the updated list.

Highlighting heavy media use and a tendency toward language not easily reduced to a set of keywords, the results displayed in Figures 2-4 suggest that a keyword approach is ill-equipped to identify young people's election discourse on Twitter.

Importantly, the presence of media and a lack of keywordifiability represent distinct challenges to a keyword approach. While young people tended to use more media, they used about the same amount of text as older generations – that is, more media does not imply less text. Despite this generational parity in text length, the actual words young people used do not appear to lend themselves to a keyword-based approach.

[1]Chi-squared tests for independence between the other cohorts did not show any statistical significance.

Critical Discourse Analysis

We concluded our investigation by critically analyzing election tweets authored by young users that a keyword approach failed to pick up on. What do we miss when we neglect these tweets, and how does it relate to issues of measurement and to our collective understanding of participatory politics?

First, we saw a politically polarized landscape take shape in the set of mislabeled tweets, with clear divisions between Biden supporters and Trump supporters. While the political opinion space cannot be reduced to two ideological camps, and a hyperfocus on polarization obscures more fundamental issues of social cohesion spurred by white supremacy (Kreiss & McGregor, 2021), the structure of US presidential elections inevitably prompts party identification and polarized behavior. Attitudinal differences emerged in the tweets, signaling distinct relationships with state power; for example, in the lead up to the election, Trump supporters expressed entitlement ("we have to win," "I want a thunderous win"), whereas Biden supporters expressed unease ("it's the United States of Anxiety," "I'm terrified for tomorrow"). Structural differences also pointed to divergent relationships with power, as a small number of prolific Trump supporters dominated the pro-Trump discourse, while support for Biden was more dispersed. Although users were easily grouped by their favored candidate, variation within those camps – particularly the Biden camp – was evident. For instance, when the race was called for Biden, many supporters celebrated Trump's loss ("#byetrump," "bye sweetie"), others – especially users of color – celebrated Harris's historic victory[2], and a few (white men) directly celebrated Biden's win.

We point out these partisan dynamics and their nuances because they provide high-level context for the examined tweets and inform the processes of in-group construction that arise. They also explain why we mostly focus on the behavior of Biden supporters in the following subsections, as this group far outnumbered the group of Trump supporters. Although there were a near equal amount of tweets from both sides, as mentioned, they were concentrated differently: 46 registered Republicans produced 187 tweets, while 151 registered Democrats produced 221 tweets. Therefore, we can speak to discursive patterns and strategies among Biden supporters without overdetermining the voices of a select few users; the same cannot be said of Trump supporters. Still, we comment on the pro-Trump tweet content where salient.

[2]Tweets celebrating Harris tended to follow a similar formula, with an image or video of Harris alongside text noting her win; as such, they seem more interested in the optics and historical nature of the event, rather than the policy implications.

Humor

We found that humor – largely conveyed through memes, visual media, and cultural references – pervaded the set of Biden supporters' tweets that were misclassified. This stood in contrast to the serious and assertive tone of Trump supporters, who embraced a matter-of-fact, pseudo-journalistic style of tweeting; for example, they rooted claims of voter fraud in "basic math and statistical probability" and posted video footage of election-related demonstrations with highly descriptive captions (e.g., "multiple protesters crowd the street as police stand by").

A great deal of the humorous pro-Biden tweets seemed to exclusively exist to mock Trump. For instance, a video of a child having a tantrum was captioned with the text, "Tr*mp's advisors explaining to him that he lost" (Figure 5, left); a similar video included the text, "actual footage of Donald Trump leaving the White House."[3] Users likewise took aim at Trump's erratic behavior during vote-counting, with one person tweeting, "pretty crazy how an NBA team can be winning at halftime, then the other team can score more points and win. I think this is a huge problem that needs to be sorted out." The mocking extended to Trump supporters more broadly, as users made fun of their rigidity (see Figure 5, right) and perception of the left ("we're all fucked up on coke at the gender neutral bathroom in antifa headquarters"). For the creators and viewers alike, these tweets served as a creative outlet for the frustration felt toward the sitting president and his base. Often, they went even deeper than frustration, providing sincere commentary on how political leadership in the US impact who and what is deemed socially acceptable (e.g., by hyperbolizing how the right views queerness); in this way, they recognized the stakes of the election, and took a serious position on the harms of Trump's reign. Finally, these tweets humiliated Trump and his power, while bonding those who were in on the joke(s) (Steele, 2016).

This bonding was further facilitated by the many tweets that found their humor through cultural references, such as the TV show *It's Always Sunny in Philadelphia* (Figure 6, left) and the video game *Among Us* (Figure 6, right). These tweets fostered in-group bonding by awarding meaning and credibility only to those familiar with the reference (Steele, 2016). They actively worked to create a shared context by providing a template for participation, and their recurrence is a testament to that. The same can be said of memes more generally.

[3]Note that neither of these tweets are keywordifiable, due to the "*" in "Tr*mp" in the former, and the fact that the text referenced in the latter is contained in the attached video rather than the body of the tweet.

Figure 5: Examples of Tweets Mocking Trump and Republicans

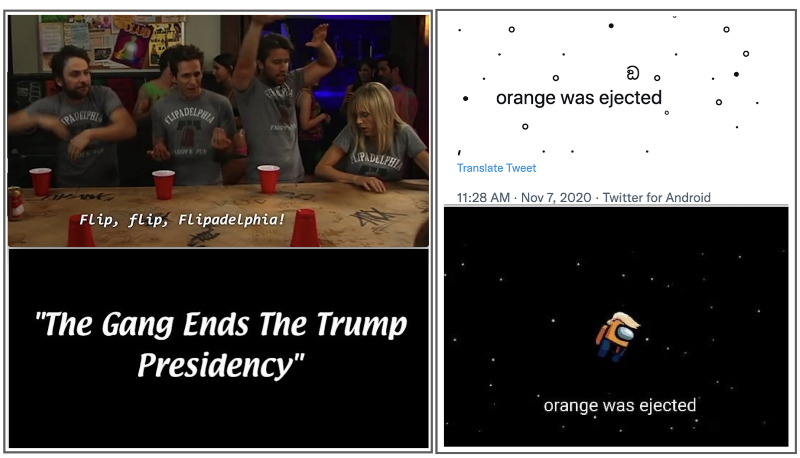

Figure 6: Examples of Tweets with Humor Drawn from Cultural References

Memes establish discursive molds that lend themselves to easy repetition and remix, which invites others into the conversation (Mina, 2019). While a number of memes showed up in our set of tweets, here we focus on the election map meme as a useful example – a selection of four instances is shown in Figure 7. The election map meme worked to (re)contextualize the election and its events by taking one of its salient symbols and repurposing it; its use demonstrates political knowledge, as one must understand something in order to remix it. In all cases, it conveyed an absurdist orientation toward the election not generally welcome in formal political spaces, whether by foregrounding non-political figures (like Britney Spears or a frog), or associating an entirely blue map with a silly, made-up platform (like producing Mama Mia 3 or enacting a new menu at Taco Bell). The map meme's absurdist elements downplay the importance of the election, acknowledging a certain powerlessness within the electoral system. Despite employing overt political symbols, the meme does not invite debate, but instead channels that powerlessness around the election into a generative joke, and in doing so, makes a claim for a different type of ownership over the election process (Christiansen & Hanson, 1996; Jenkins et al., 2016). Meanwhile, the (re)contextualization separates the meme from mainstream discourse somewhat, maintaining the claim of ownership for the in-group alone, and a non-threatening appearance in the dominant group's eyes (Steele, 2016).

Figure 7: Four instances of the election map meme. Note that the tweets that look like they were authored by Biden were instead photoshopped for the sake of the joke.

The humorous tweets discussed here embody participatory politics through their playful, interactive nature, and the way they leveraged cultural references to bond users and reconfigure political power. Yet, they present clear obstacles for keyword-based detection, as much of their meaning is tied to media (which also holds for the reporter-esque pro-Trump tweets mentioned earlier), external cultural referents, and/or intentionally obscured symbols.

Collective Emotion

Perhaps more than anything, emotions were centered in our set of examined tweets. As a whole, the tweets depicted "election Twitter" as a space to cultivate mood and circulate emotion, which stands in stark contrast to the way emotions are typically devalued in favor of logic and rationality in formal political settings (Ahmed, 2014; Blumler & Kavanagh, 1999; Steele, 2016). As with humor, this feature primarily resided in the Biden camp; although Trump supporters clearly expressed anger toward Democrats and left-leaning media, they mainly used that emotion to proffer intellectual claims rather than communicate feeling.

The central role of affect was largely evidenced by the general orientation toward mainstream media among Biden supporters: These users cited media reporting, and waited on news organizations to call the race. In these ways, and unlike many of the Trump supporters, they did not present as journalists or analysists themselves, but something else, primarily engaged with the *experience* of the election. This emphasis on subjective experience echoes the way bloggers have been described, which has contributed to their marginalization as writers (Papacharissi & Meraz, 2012). However, it should not be taken as a weakness, as it aims the spotlight at those actually participating in the discourse and encourages them to bring their full selves to the conversation. After all, electoral outcomes (and politics more generally) matter deeply at the level of the individual.

There were many moments categorized by an outpouring of emotion, such as anxiety in the lead-up to the results (as mentioned), and joy after the race was called. In these cases, users often included images and videos capturing rich emotional expression (see Figure 8), which served to heighten affective signals. Emojis, gifs, and other built-in Twitter features were used for similar ends. Adding to this, an entire genre revolving around calls for self-care and "sanitizing the timeline" cropped up. The calls for self-care acknowledged the election's emotional intensity, and directed readers to attend to how they feel ("how are you treating yourself today? you deserve it"). Relatedly, the attempts to "sanitize the timeline" mostly involved cute animal pictures meant to distract viewers from the election and its emotional toll. Overall, these tweets emphasized *feeling* to their audiences, and positioned it as a critical aspect of the election; moreover, they actively intervened on feelings of anxiety and overwhelm, creating a community of care and signaling a sense of political self-efficacy.

Participatory politics is evident in these appeals to emotion, as they renegotiate what matters during the election and who gets decide. Through

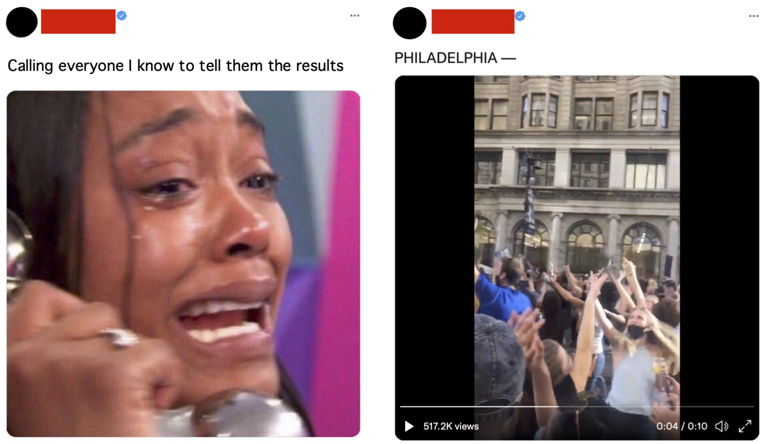

Figure 8: Examples of Tweets That Use Media to Convey Emotion

that process of renegotiation, they also forge social ties and perform acts of solidarity that communicate political agency/empowerment. Like the humorous tweets, the affective tweets too resist a keyword approach for being media-rich and/or deliberately averting straightforward election speech in the name of community care.

As a whole, the false negative election tweets – whether emotional, humorous, partisan, or otherwise – highlight how an automated keyword approach misses important examples of political speech, categorizing them instead as irrelevant and functionally erasing them from the election discourse.

Discussion

We found that a keyword-based approach to data collection is not well-suited to detect young people's participation in online election discourse. First, young people communicated about the election in large part through media, which is beyond the scope of a keyword method. Second, many of their election tweets seemed to be absent of standard election language, which violates the fundamental assumption of the approach. Our critical discourse analysis illuminated each of these challenges, and uncovered how they specifically disregard manifestations of participatory politics. It is likely, too, that the problem will only exacerbate as "algospeak" – purposely obscured language to evade algorithmic content moderation – becomes more widespread (Lorenz, 2022).

The limits of a keyword approach are closely linked to the enduring narrative of youth political apathy. In the shadow of this narrative, incomplete data relating to young people's political participation could easily read as evidence of political apathy. Researchers may be unlikely to realize that content is being missed, as this evidence is in line with a prominent argument. What's more, they are poised to reinforce that argument. Charging young people with political apathy is completely unhelpful. If apathy were indeed the case, publicly maligning young people for it would be a shallow and unserious response. Further, the charge makes no attempt to recognize – and thereby effectively erases – the political contributions they do make.

The participatory politics literature makes clear that young people are meaningfully engaged with politics, just in ways that are less tied to formal structures. This was underscored by our critical discourse analysis of young people's mislabeled election speech. We found that the overlooked tweets contained vivid conversations about the election. In leveraging humor and affect – itself a powerful rejection of elitist political norms – they signaled political understanding and self-efficacy, facilitated in-group bonding, and renegotiated political power. The fact that they encompassed non-standard ways of discussing the election – through absurdist memes or emotional acknowledgment, for instance – does not take away from their relevance or value. Rather, it adds to it by creating more space for critical engagement with the election that puts voters and their experiences at the center.

Therefore, future work should focus on how to better detect the political speech associated with young people. While a standard keyword approach is bound to miss a great deal, our results suggest that incorporating optical character recognition (OCR) (Chaudhuri et al., 2017) – that is, searching for linguistic signals within the attached media – would already be a major improvement. Tools for implementing OCR already exist, such as the open source Tesseract engine[4] (Amalia et al., 2018), and researchers should strive to make them more accessible and mainstream.

Beyond that, researchers should work toward developing methods that can recognize vague or nonstandard speech – i.e., speech that is *not* keywordifiable. This may involve leaning on contextual clues, such as the timestamp of a post, the language used in the replies, or the poster's history. The latter has shown promise in the case of hate speech detection (Qian et al., 2018). The concept of keywordifiability underscores the plurality of ways one can communicate about a given topic by measuring the extent to which some discourse can be reduced to a set of keywords. This is especially useful

[4]https://github.com/tesseract-ocr/tesseract/

in the context of communication on social media, where expression is free of journalistic standards and other limits; if we resort to imposing limits in our attempts to capture this expression – for instance, by decreeing that certain keywords be present – we miss much of what makes social media discourse distinctive, vibrant, and inclusive. Keywordifiability is thus helpful for determining the sufficiency of using keywords to study a particular online discourse, as well as for thinking through who and what might be discounted.

The policing of what counts as valid political participation often works to privilege the already privileged and marginalize the already marginalized. For better or for worse, researchers hold significant power in this validation process. This study shines a light on how that power may be as subtle as constructing a keyword list: The seemingly small assumption that election speech contains mainstream election text can introduce a whole suite of biases that affect how different groups' participation is understood. Here, we focused on how that impacts young people; however, our conclusions carry implications for other social groups at the margins of mainstream discourse (Blodgett et al., 2016; Fraser, 1990), which future studies should dig deeper into.

Acknowledgements

We are beyond grateful to Hana Chalmers and Alex Madaras, undergraduate research assistants, who played key roles in the qualitative coding portion of this work. We are also grateful to David Lazer for providing helpful feedback throughout the project.

Funding

S.M. was supported by the John S. and James L. Knight Foundation through a grant to the Institute for Data, Democracy & Politics at The George Washington University.

References

Ahmed, S. (2014). Introduction: Feel Your Way. In *The Cultural Politics of Emotion* (NED - New edition, 2, pp. 1–19). Edinburgh University Press. Retrieved November 30, 2021, from https://www.jstor.org/stable/10.3366/j.ctt1g09x4q.5

Amalia, A., Sharif, A., Haisar, F., Gunawan, D., & Nasution, B. B. (2018). Meme Opinion Categorization by Using Optical Character Recognition (OCR) and Naïve Bayes Algorithm. *2018 Third International Conference on Informatics and Computing (ICIC)*, 1–5. https://doi.org/10.1109/IAC.2018.8780410

Bennett, W. L., Wells, C., & Freelon, D. (2011). Communicating Civic Engagement: Contrasting Models of Citizenship in the Youth Web Sphere. *Journal of Communication*, *61*(5), 835–856. https://doi.org/10.1111/j.1460-2466.2011.01588.x

Blodgett, S. L., Green, L., & O'Connor, B. (2016). Demographic Dialectal Variation in Social Media: A Case Study of African-American English [arXiv: 1608.08868]. *arXiv:1608.08868 [cs]*. Retrieved May 4, 2022, from http://arxiv.org/abs/1608.08868
Comment: To be published in EMNLP 2016, 15 pages

Blumler, J. G., & Kavanagh, D. (1999). The Third Age of Political Communication: Influences and Features. *Political Communication*, *16*(3), 209–230. https://doi.org/10.1080/105846099198596

Center, P. R. (2019). *National Politics on Twitter: Small Share of U.S. Adults Produce Majority of Tweets* (tech. rep.). Pew Research Center. https://www.pewresearch.org/politics/2019/10/23/national-politics-on-twitter-small-share-of-u-s-adults-produce-majority-of-tweets/

Center, P. R. (2021). *Social Media Fact Sheet* (tech. rep.). Pew Research Center. Retrieved November 29, 2021, from https://www.pewresearch.org/internet/fact-sheet/social-media/

Center, P. R. (2022). *American's Views of Government: Decades of Distrust, Enduring Support for Its Role* (tech. rep.). Pew Research Center. https://www.pewresearch.org/politics/2022/06/06/americans-views-of-government-decades-of-distrust-enduring-support-for-its-role/

Chaudhuri, A., Mandaviya, K., Badelia, P., & Ghosh, S. K. (2017). Optical Character Recognition Systems. In A. Chaudhuri, K. Mandaviya, P. Badelia, & S. K Ghosh (Eds.), *Optical Character Recognition Systems for Different Languages with Soft Computing* (pp. 9–41). Springer International Publishing. https://doi.org/10.1007/978-3-319-50252-6_2

Chen, E., Deb, A., & Ferrara, E. (2022). #Election2020: The first public Twitter dataset on the 2020 US Presidential election. *Journal of Computational Social Science*, *5*(1), 1–18. https://doi.org/10.1007/s42001-021-00117-9

Christiansen, A. E., & Hanson, J. J. (1996). Comedy as cure for tragedy: Act up and the rhetoric of aids. *Quarterly Journal of Speech*, *82*(2), 157–170. https://doi.org/10.1080/00335639609384148

Cohen, C. J., & Kahne, J. (2012). *Participatory Politics: New Media and Youth Political Action* (tech. rep.). Youth and Participatory Politics Research Network.

Driscoll, K., & Walker, S. (2014). Big Data, Big Questions| Working Within a Black Box: Transparency in the Collection and Production of Big Twitter Data [Number: 0]. *International Journal of Communication*, *8*(0), 20. Retrieved July 12, 2022, from https://ijoc.org/index.php/ijoc/article/view/2171

Fairclough, N. (2013). *Critical Discourse Analysis: The Critical Study of Language* (2nd ed.). Routledge. https://doi.org/10.4324/9781315834368

Flanagan, C., & Levine, P. (2010). Civic Engagement and the Transition to Adulthood [Publisher: Princeton University]. *The Future of Children*, *20*(1), 159–179. Retrieved July 12, 2022, from https://www.jstor.org/stable/27795064

Fraser, N. (1990). Rethinking the Public Sphere: A Contribution to the Critique of Actually Existing Democracy [Publisher: Duke University Press]. *Social Text*, (25/26), 56–80. https://doi.org/10.2307/466240

Freelon, D., Lopez, L., Clark, M., & Jackson, S. (2018). *How Black Twitter and Other Social Media Communities Interact With Mainstream News* (tech. rep.). Knight Foundation. https://knightfoundation.org/features/twittermedia/

Gadde, V., & Beykpour, K. (2020). Additional steps we're taking ahead of the 2020 US Election. https://blog.twitter.com/en_us/topics/company/2020/2020-election-changes

Gallagher, R. J., Doroshenko, L., Shugars, S., Lazer, D., & Foucault Welles, B. (2021). Sustained Online Amplification of COVID-19 Elites in the United States [Publisher: SAGE Publications Ltd]. *Social Media + Society*, *7*(2), 20563051211024957. https://doi.org/10.1177/20563051211024957

Gorwa, R., & Guilbeault, D. (2020). Unpacking the Social Media Bot: A Typology to Guide Research and Policy. *Policy & Internet*, *12*(2), 225–248. https://doi.org/10.1002/poi3.184

Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., & Lazer, D. (2019). Fake news on Twitter during the 2016 U.S. presidential election [Publisher: American Association for the Advancement of Science]. *Science*, *363*(6425), 374–378. https://doi.org/10.1126/science.aau2706

Hughes, A. G., McCabe, S. D., Hobbs, W. R., Remy, E., Shah, S., & Lazer, D. M. J. (2021). Using Administrative Records and Survey Data to Construct Samples of Tweeters and Tweets. *Public Opinion Quarterly*, *85*(S1), 323–346. https://doi.org/10.1093/poq/nfab020

Jackman, S., & Spahn, B. (2021). Politically Invisible in America [Publisher: Cambridge University Press]. *PS: Political Science & Politics*, *54*(4), 623–629. https://doi.org/10.1017/S1049096521000639

Jackson, S. J., Bailey, M., & Welles, B. F. (2020). *#HashtagActivism: Networks of Race and Gender Justice*. MIT Press.

Jenkins, H. (2006). *Fans, Bloggers, and Gamers: Exploring Participatory Culture*. NYU Press.

Jenkins, H., Billard, T. J., Close, S., Elsayed, Y., Forelle, M. C., Lopez, R., & Yang, E. (2017). Participatory Politics [Num Pages: 16]. In *Keywords in Remix Studies*. Routledge.

Jenkins, H., Ito, M., & boyd danah, d. (2015). *Participatory Culture in a Networked Era: A Conversation on Youth, Learning, Commerce, and Politics* (1st ed.). Polity.

Jenkins, H., Shresthova, S., Gamber-Thompson, L., Kligler-Vilenchik, N., & Zimmerman, A. M. (2016). *By Any Media Necessary: The New Youth Activism*. New York University Press.

Kennedy, C., & Deane, C. (2017). *A basic question when reading a poll: Does it include or exclude nonvoters?* (Tech. rep.). Pew Research Center. https://www.

pewresearch.org/fact-tank/2017/02/16/does-poll-include-or-exclude-nonvoters/
Khosla, V. (2022). Political Partisanship, Extreme Polarization and Youth Voter Turnout in 2020. *Penn Journal of Philosophy, Politics & Economics*, *17*(1). https://repository.upenn.edu/spice/vol17/iss1/5
King, G., Lam, P., & Roberts, M. E. (2017). Computer-Assisted Keyword and Document Set Discovery from Unstructured Text. *American Journal of Political Science*, *61*(4), 971–988. https://doi.org/10.1111/ajps.12291
King, M., & Weisman, J. (2022). Young Voters Are Fed Up With Their (Much) Older Leaders. *The New York Times*. Retrieved July 14, 2022, from https://www.nytimes.com/2022/07/14/us/politics/youth-voters-midterms-polling.html
Kligler-Vilenchik, N., & Literat, I. (2018). Distributed Creativity as Political Expression: Youth Responses to the 2016 U.S. Presidential Election in Online Affinity Networks. *Journal of Communication*, *68*(1), 75–97. https://doi.org/10.1093/joc/jqx005
Kligler-Vilenchik, N., & Literat, I. (2020). Youth Digital Participation: Now More than Ever. *Media and Communication*, *8*(2), 171–174. https://doi.org/10.17645/mac.v8i2.3180
Kreiss, D., & McGregor, S. (2021). Polarization Isn't America's Biggest Problem—or Facebook's. *Wired*. Retrieved November 30, 2021, from https://www.wired.com/story/polarization-isnt-americas-biggest-problem-or-facebooks/
Lane, D. S. (2020). Social media design for youth political expression: Testing the roles of identifiability and geo-boundedness [Publisher: SAGE Publications]. *New Media & Society*, *22*(8), 1394–1413. https://doi.org/10.1177/1461444819879103
Lane, D. S., Thorson, K., & Xu, Y. (2023). Uninterested and unequal?: Examining SES-based gaps in youth political behavior on social media. *Information, Communication & Society*, *26*(4), 663–681. https://doi.org/10.1080/1369118X.2021.1968923
Literat, I., & Kligler-Vilenchik, N. (2019). Youth collective political expression on social media: The role of affordances and memetic dimensions for voicing political views [Publisher: SAGE Publications]. *New Media & Society*, *21*(9), 1988–2009. https://doi.org/10.1177/1461444819837571
Literat, I., & Kligler-Vilenchik, N. (2021). How Popular Culture Prompts Youth Collective Political Expression and Cross-Cutting Political Talk on Social Media: A Cross-Platform Analysis [Publisher: SAGE Publications Ltd]. *Social Media + Society*, *7*(2), 20563051211008821. https://doi.org/10.1177/20563051211008821
Longley, R. (2022). Requirements to Become President of the United States [Section: ThoughtCo]. Retrieved July 13, 2022, from https://www.thoughtco.com/requirements-to-serve-as-president-3322199
Lorenz, T. (2022). Internet 'algospeak' is changing our language in real time, from 'nip nops' to 'le dollar bean'. *Washington Post*. Retrieved August 8, 2022, from https://www.washingtonpost.com/technology/2022/04/08/algospeak-tiktok-le-dollar-bean/
Lotan, G., Graeff, E., Ananny, M., Gaffney, D., Pearce, I., & Boyd, D. (2011). The Arab Spring| The Revolutions Were Tweeted: Information Flows during the 2011

Tunisian and Egyptian Revolutions [Number: 0]. *International Journal of Communication*, *5*(0), 31. Retrieved August 8, 2022, from https://ijoc.org/index.php/ijoc/article/view/1246

Mejia, R., Beckermann, K., & Sullivan, C. (2018). White lies: A racial history of the (post)truth. *Communication and Critical/Cultural Studies*, *15*(2), 109–126. https://doi.org/10.1080/14791420.2018.1456668

Mina, A. X. (2019). *Memes to Movements: How the World's Most Viral Media Is Changing Social Protest and Power* [Google-Books-ID: iMNUDwAAQBAJ]. Beacon Press.

Mitchell, A., Shearer, E., & Stocking, G. (2021). News on Twitter: Consumed by Most Users and Trusted by Many. Retrieved August 16, 2022, from https://www.pewresearch.org/journalism/2021/11/15/news-on-twitter-consumed-by-most-users-and-trusted-by-many/

Munger, K. (2017). Tweetment Effects on the Tweeted: Experimentally Reducing Racist Harassment. *Political Behavior*, *39*(3), 629–649. https://doi.org/10.1007/s11109-016-9373-5

Nwanevu, O. (2014). The Right to Run [Section: Politics]. *Slate*. Retrieved July 13, 2022, from https://slate.com/news-and-politics/2014/10/age-of-candidacy-laws-should-be-abolished-why-18-year-olds-should-be-able-to-run-for-public-office.html

O'Neill, A. (2022). *Voter turnout rates* among selected age groups in U.S. presidential elections from 1964 to 2020* (tech. rep.). Statista. https://www.statista.com/statistics/1096299/voter-turnout-presidential-elections-by-age-historical/

Papacharissi, Z., & Meraz, S. (2012). Blogging Culture. In *The International Encyclopedia of Media Studies*. John Wiley & Sons, Ltd. https://doi.org/10.1002/9781444361506.wbiems074

Penney, J. (2019). "It's My Duty to Be Like 'This Is Wrong'": Youth Political Social Media Practices in the Trump Era". *Journal of Computer-Mediated Communication*, *24*(6), 319–334. https://doi.org/10.1093/jcmc/zmz017

Penney, J. (2020). 'It's So Hard Not to be Funny in This Situation': Memes and Humor in U.S. Youth Online Political Expression [Publisher: SAGE Publications]. *Television & New Media*, *21*(8), 791–806. https://doi.org/10.1177/1527476419886068

Poll, H. Y. (2022). *Harvard Youth Poll: Top Trends and Takeaways* (tech. rep.). Harvard Kennedy School Institute of Politics. https://iop.harvard.edu/youth-poll/spring-2022-harvard-youth-poll

Prensky, M. (2001). Digital Natives, Digital Immigrants Part 1 [Publisher: MCB UP Ltd]. *On the Horizon*, *9*(5), 1–6. https://doi.org/10.1108/10748120110424816

Putnam, R. D. (2000). *Bowling Alone: The Collapse and Revival of American Community*. Simon & Schuster.

Qian, J., ElSherief, M., Belding, E., & Wang, W. Y. (2018). Leveraging Intra-User and Inter-User Representation Learning for Automated Hate Speech Detection. *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, 118–123. https://doi.org/10.18653/v1/N18-2019

Shugars, S., Gitomer, A., McCabe, S., Gallagher, R. J., Joseph, K., Grinberg, N., Doroshenko, L., Welles, B. F., & Lazer, D. (2021). Pandemics, Protests, and Publics: Demographic Activity and Engagement on Twitter in 2020. *Journal of Quantitative Description: Digital Media*, *1*. https://doi.org/10.51685/jqd.2021.002

Sloam, J., & Henn, M. (2019). Rejuvenating Politics: Young Political Participation in a Changing World. In J. Sloam & M. Henn (Eds.), *Youthquake 2017: The Rise of Young Cosmopolitans in Britain* (pp. 17–42). Springer International Publishing. https://doi.org/10.1007/978-3-319-97469-9_2

Steele, C. K. (2016). The Digital Barbershop: Blogs and Online Oral Culture Within the African American Community [Publisher: SAGE Publications Ltd]. *Social Media + Society*, *2*(4), 2056305116683205. https://doi.org/10.1177/2056305116683205

Tufekci, Z. (2017). *Twitter and Tear Gas: The Power and Fragility of Networked Protest* [Google-Books-ID: x7S_DgAAQBAJ]. Yale University Press.

Van Deth, J. W. (2016). What Is Political Participation? [ISBN: 9780190228637]. https://doi.org/10.1093/acrefore/9780190228637.013.68

Varol, O., Ferrara, E., Davis, C. A., Menczer, F., & Flammini, A. (2017). *Online Human-Bot Interactions: Detection, Estimation, and Characterization* (tech. rep.) [Publication Title: arXiv e-prints ADS Bibcode: 2017arXiv170303107V Type: article]. Retrieved August 8, 2022, from https://ui.adsabs.harvard.edu/abs/2017arXiv170303107V

Vromen, A., Loader, B. D., Xenos, M. A., & Bailo, F. (2016). Everyday Making through Facebook Engagement: Young Citizens' Political Interactions in Australia, the United Kingdom and the United States [Publisher: SAGE Publications Ltd]. *Political Studies*, *64*(3), 513–533. https://doi.org/10.1177/0032321715614012

Wattenberg, M. P. (2015). *Is Voting for Young People?* (3rd ed.). Routledge. https://doi.org/10.4324/9781315663746

White, A., & Nguyen, A. (2022). How Often Do People Vote While Incarcerated? Evidence from Maine and Vermont [Publisher: The University of Chicago Press]. *The Journal of Politics*, *84*(1), 568–572. https://doi.org/10.1086/714927

Zhu, A. Y. F., Chan, A. L. S., & Chou, K. L. (2019). Creative social media use and political participation in young people: The moderation and mediation role of online political expression. *Journal of Adolescence*, *77*(1), 108–117. https://doi.org/10.1016/j.adolescence.2019.10.010

Appendix: Keyword List

election
presdebate
vpdebate
democratic
gop
dnc
rnc
politics
political
voter
senate
senator
2020election
election2020
vote
voting
electionday
votebymail
votersuppression
ballot
mailin
mail-in
mail in
russiahoax
qanon
obamagate
mailinballots
nakedballots
presidential
vote-by-mail
votingsquad
votethemout
wewillvote
mike pence
michael pence
mikepence
michaelpence
pence
kamala harris
kamala
harris
spike cohen
angela walker
kamalaharris
senkamalaharris
joe biden
joebiden
biden
votebiden
bluewave2020
votebidenharris2020
ridinwithbiden
nomalarkey
biden2020
bidenharris2020
bidenforpresident
bidenkamala2020
joebiden2020
bidenharris
votehimout
dumptrump
nevertrump
bluewave
fucktrump
bidenwarroom
voteblue
demconvention
votebluetosaveamerica
wakeupamerica
trumpisanationaldisgrace
trumpvirus
trumpisalaughingstock
traitortrump
jo jorgensen
jojorgensen
joanne marie jorgensen

jorgensen2020
beboldvotegold
donald trump
don trump
realdonaldtrump
donaldtrump
donaldjtrump
donald j trump
trump
trumpwarroom
teamtrump
the donald
trump2020
maga
draintheswamp
keepamericagreat
neverbiden
trumppence2020
makeamericagreatagain
kag
presidenttrump
notmypresident
americafirst
redwave
votered
sleepy joe
sleepyjoe
hidenbiden
creepyjoebiden
bidenukrainescandal
rnc2020
kag2020
maga2020
trump2020landslide
tulsatrumprally
voteredtosaveamerica
trumpforpresident
backtheblue
howiehawkins
howie hawkins
howiehawkins2020
hawkins2020
blackvotesmatter

Appendix: Codebook

DESCRIPTION OF TASK:

Our task is to identify tweets related to the 2020 U.S. presidential election.

PRE-FILLED FIELDS:

Title	**Description/Definition**
LINK	Link to tweet (str)
USERID	User ID of tweet's author (int)
AGE	Age of tweet's author, as of 6/2021 (float)
TEXT	Full text of tweet (str)
FAVORITE COUNT	Favorite count of tweet, as of 12/2021 (int)
RETWEET COUNT	Retweet count of tweet, as of 12/2021 (int)
CREATED AT	Date and time at which tweet was posted, EST (str)
MEDIA	"True" if tweet includes media, "False" otherwise (bool)
MEDIA COUNT	Number of media attachments in tweet (int)
MEDIA TYPE	Set of strings describing type of media attachments [photo/video/animated gif] (set/list)

FIELDS TO CODE:

THE FOLLOWING VARIABLES ARE NOT MUTUALLY EXCLUSIVE!

Title	Description/Definition
ELECTION P	Coded as 1 if the tweet is about the PRESIDENTIAL election, and 0 otherwise. The tweet must strongly reference the 2020 presidential election and its associated events. This includes tweets that mention one or more of the (vice) presidential candidates (the central candidates were Donald Trump, Joe Biden, Howie Hawkins, and Jo Jorgensen; a more complete list of candidates can be found here). This also includes tweets that reference vote counting, voting methods, election-related activities (e.g., rallies), voting fraud, etc., in ways that firmly tie them to the events of the 2020 presidential election (e.g., calls to stop vote counting). Mention of battleground states or states that were slow to be called may also heavily (but not necessarily) indicate that the tweet is related to the presidential election. Be inclusive when marking this category.
ELECTION V	Coded as 1 if the tweet seems to be about the presidential election, but is vague, absent of obvious keywords, or otherwise ambiguous; 0 otherwise. This can include: vague panicking about the results of the election, vague wondering about the current state of the election, vague confusion surrounding election results, vague concern for what could happen after the election, etc. Includes tweets about voter turnout and/or efforts to suppress it. Put differently, in order to be meaningful, the tweet seems to rely on the notion that most people reading it are primed in their thinking by the election and its events, which were particularly chaotic and uncertain in 2020. Use the replies and the date to help guide your intuition. Be inclusive when marking this category.

Appendix

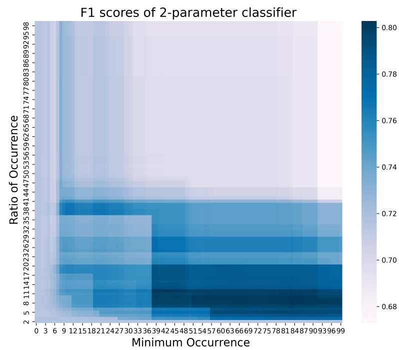

Figure 1: Election tweet classifier F1 scores for minimum occurrence parameter α between 0 and 99 along the x-axis (i.e., a keyword must appear that many times in the training set's election tweets), and ratio of occurrence parameter β between 0 and 99 along the y-axis (i.e., the ratio of occurrence for a keyword between election tweets and non-election tweets in the training set must be at or above that threshold). The highest F1 score achieved is 80.3%, which occurs when $\alpha = 83$. This is in line with the performance of our classifier based on an expert-curated keyword list, which achieved an F1 score of 80.9%.