Using Twitter Data to Predict Violent Events

Honor's Political Science Thesis

Kelsea Hull

2018

Intro

With the popularity of social media, platforms such as Twitter have become the perfect place for extremist groups to espouse ideas and recruit members. The goal of this project was to determine if the frequency of use of twitter by extremist groups or individuals associated with these groups and ideologies reliably correlates to the lead up to a violent event, such as the shooting perpetrated by Dylann Roof in Charleston.

After this shooting, evidence of the perpetrators activity online surfaced. He had a blog called 'Last Rhodesian' where he openly talked about his desire to remove black people form his home state. The week directly preceding his attack, he posted a last message about his frustration of other online extremist's unwillingness to act, and how he had decided to take it upon himself and hoped to inspire others to do the same.[1] After this came to light, the effect of social media on radicalization within the US became a more popular topic, and it was this that ultimately inspired this project.

Being able to use readily available information such as twitter data, or other social media platforms, to predict a violent event would give law enforcement a greater chance of preventing the event. Not only would lives potentially be saved if a violent event were prevented, but extremist ideology driven events have the to potential to inspire more of the same, or to inspire others to adopt the same beliefs. Preventing these events from occurring could then prevent the spread of motivation to commit such acts. This project would then hope to determine if a pattern can be found, and whether that pattern could be used to predict future events.

---

[1] Dylann Roof, "Manifesto," *Last Rhodesian* (blog), June 20, 2015, accessed May 10, 2018, https://www.documentcloud.org/documents/2108059-lastrhodesian-manifesto.html.

Hypothesis

The goal of this project is to analyze the Twitter data from three months prior to a far-right wing ideology driven violent event for the frequency of certain keywords, and then compare this to the total Twitter data from my larger sample between March 1, 2017 and January 2, 2018. The hope is to find a pattern that can be used to predict future events. In the discussion surrounding Roof's trial, it was released that he traced his own radicalization back to when he searched 'black on white crime' in Google. He found the Council of Conservative Citizens (CofCC) website, which is a known white supremacist group, and their Twitter was one of the accounts I attempted to gather data from. After Charleston, there was a good deal of discussion about online radicalization of right wing extremist youth in the United States and the proliferation of dangerous ideologies such as violent white supremacy on the internet.[2] This led to the idea of tracking social media activity leading up to a past attack, and this research is what mostly inspired the accounts and keywords that were used.

Currently, research is being done on using digital sources such as social media and online forums to detect radicalization, but it is mainly focusing on identifying potential lone wolf actors before terror events. One idea proposed is to use sentiment analysis and author recognition on social media postings to create a semi-automated tool for identifying lone wolves.[3] The research presented by Cohen et al. focuses on detecting three specific weak signs of violent lone wolf behavior, namely leakage, fixation, and identification. In this context: leakage is the tendency of lone wolves to tell others of their plans before carrying them out; fixation is when the lone wolf

---

[2] Rebecca Hersher, "What Happened When Dylann Roof Asked Google For Information About Race?," *NPR*, Jan 10, 2017, accessed May 10, 2018, https://www.npr.org/sections/thetwo-way/2017/01/10/508363607/what-happened-when-dylann-roof-asked-google-for-information-about-race.

[3] Katie Cohen et al., "Detecting Linguistic Markers for Radical Violence in Social Media," Terrorism and Political Violence 26, no. 1 (undefined): 246-56, http://dx.doi.org/10.1080/09546553.2014.849948.

consistently fixates on one group, and frequently goes had in hand with in-group/out-group rhetoric; identification is the desire to relate oneself to military/police weapons, previous attackers, or warrior attitudes. A similar approach to using digital writings on social media and forums to predict lone wolf radicalization uses hyperlink analysis and natural language processing in methods similar to those used to identify child pornography networks online.[4] The goal of the research conducted by Brynielsson et al. is again to detect weak signs of lone wolf violence before it happens to prevent individuals from acting upon them. The analysis I am attempting differs in what it is hoping to predict. Rather than trying to identify individuals who have been radicalized to the point of lone wolf violence, this project hopes to analyze social media content for to predict the occurrence of events. The hypothesis for this project is that there will be a higher frequency of keywords leading up to an event as compared to other times. The first step in gathering data was deciding which accounts to collect from on twitter.

Data Description

The initial list of accounts to look through came from the Southern Poverty Law Center (SPLC) Extremist Files page.[5] I took the accounts of leaders of right wing extremist groups and compiled a list of about 100, drawing from three different categories listed on the SPLC page: anti-immigrant, alt-right, and white nationalist. While this was not an exhaustive list, it was kept short for the purpose of ease of analysis for the length and scope of this project. These group accounts were also taken from the reports surrounding the trial of Dylann Roof. From there, the problem of not being able to view accounts arose, as Twitter has recently started banning accounts that violate Twitter community standards, including those that are guilty of excessively

[4]Joel Brynielsson et al., "Harvesting and analysis of weak signals for detecting lone wolf terrorists," Security Informatics 2, no. 1 (undefined): 11, http://dx.doi.org/10.1186/2190-8532-2-11.

[5] "Extremist Files," SPLC Southern Poverty Law Center, accessed May 10, 2018, https://www.splcenter.org/fighting-hate/extremist-files.

violent hate speech.[6] Twitter banning accounts also creates the problem of historical data not being visible to the public, or collectible for research. At the start of data collection, I had decided that if I could find 40 accounts worth of data, this would be a large enough sample for this project. After having the initial list of 100 cut down significantly, I went through followers of the accounts on my list that I could access to see those who frequently retweeted their posts and chose accounts from there to get a total of 42 accounts to collect tweets from.[7] The process of choosing which accounts to collect from was a little bit random, but did involve looking at how often they appeared in the mentions of the initial accounts, and how many followers they had.

I gathered twitter data through the Twitter API, which only allows the collection of the 3200 most recent tweets from any account.[8] All the actual programming has been done through Jupyter Notebooks in python. To use the Twitter API, it is necessary to create an account and get a unique access key. Then, I had to collect the data from each account separately. I used a publicly available GitHub repository to collect the twitter data.[9] While this was the easiest way to get data, it has some problems. I was not able to obtain all tweets after a certain beginning date: users may tweet so frequently that the 3200 tweets do not stretch back further than a few months. This means for some accounts I have data going back almost to the beginning of Twitter's popularity, while others go back only a few years from the request date of Jan 2, 2018. The average length of time that I have collected tweets from for each account is 501 days. This can

---

[6] Emily Sullivan, "Twitter Says It Will Ban Threatening Accounts, Starting Today," *NPR*, Dec. 18, 2017, accessed May 10, 2018, https://www.npr.org/sections/thetwo-way/2017/12/18/571622652/twitter-says-it-will-ban-threatening-accounts-starting-today.

[7] Reference table 1 for a complete list of accounts.

[8] "API Reference," Twitter Developer, 2018, accessed May 10, 2018, https://api.twitter.com/1.1/statuses/user_timeline.json.

[9] Yanofsky, David, "tweet_dumper.py," (2013), GitHub repository, accessed January 2018, https://gist.github.com/yanofsky/5436496.

be problematic when using the data as samples to run tests on and will be discussed later. Each account's twitter data created a separate .csv file which then needed to be combined into one file, so it could be useful. In total I have 76034 tweets from 42 accounts.

The next step with the Twitter data was to choose keywords to track and then count the number of times each word appeared per tweet. The list of words I chose was again inspired by the SPLC Extremist Files page but was also taken from research on the event that originally inspired this project. During his trial, Roof's manifesto and prison journal were released, which gave details of the kind of language frequently used by extremist white supremacists online, and this provided a good base for keywords.[10] Some were also taken from combing through the accounts or blogs themselves, as well as the social media site 'Gab', which was created by far-right users who were banned from Twitter for hate speech.[11] For each keyword, the number of times used per tweet was counted, as well as the mean number of times used per month for the total data, and the mean number of times each word was used per day for the total data. The means were then normalized by the number of active users that day, which I defined as the number of users who had tweeted in the 14 days leading up to that day.[12] This mean per day count is what will be used for the samples for the difference of means test.[13] This list of keywords is not meant to be exhaustive, but was just a small sample of buzzwords taken from sources such as Dylann Roof's blog and the press surrounding Charlottesville.

Analysis

[10] "The Dylann Roof Trial: The Evidence," *New York Times*, Dec 9, 2016, accessed May 10, 2018, https://www.nytimes.com/interactive/2016/12/09/us/dylann-roof-evidence.html.

[11] "The Free Speech Social Network," Gab, 2018, accessed May 10, 2018, https://gab.ai/.

[12] Reference Graph 3.

[13] Reference table 1 for a complete list of keywords.

After collecting all the data, it was necessary to find the proper sample sizes to run tests on. While I was originally intending to use the Charleston shooting, this was not feasible since I did not have tweets from a reasonable number of active accounts in 2015. I then chose the death of a protestor in Charlottesville on August 12, 2017 as my event to analyze. This was an incident where there was a scheduled white nationalist rally, and people showed up to counter-protest. One of the rally participants drove into a group of counter-protestors, killing one and injuring at least 19 others.[14] After selecting an event, my smaller sample includes all tweets beginning three months before this event. The larger sample was a little harder to put together, since it was difficult to know where to start. To choose a starting point, I graphed the total number of users that tweeted each day, and chose to start March 1, 2017, and end on January 2, 2018 when the data ends, excluding the three months that make up the smaller sample.[15]

One interesting thing found just by chance through this data manipulation was the potential for twitter data to reveal events looking at it historically. While manipulating the data one of the things I tried was to graph the total number of tweets per day or the full twitter data, and this showed a strange spike surrounding Halloween 2015. In the months before and after, there was no more than 10 tweets per day, but in the week surrounding Halloween the total jumped up to 300, which seemed unlikely.[16] Through this, and through googling the data I found that the spike in tweets was related to a white nationalist rally focused on securing a white ethno-

[14] Sheryl Gay Stolberg and Brian M. Rosenthal, "Man Charged After White Nationalist Rally in Charlottesville Ends in Deadly Violence," *New York Times*, Aug. 12, 2017, accessed May 10, 2018, https://www.nytimes.com/2017/08/12/us/charlottesville-protest-white-nationalist.html.

[15] Reference Graph 1.

[16] Reference Graph 4.

state.[17] This is just one example of how social media analysis such as this could be used to predict future events, or to study trends in the past.

To emphasize the utility of this type of data in studying past events, I have included the graph of the number of tweets per day normalized by the number of active users, and the graph of the number of users per day normalized by the number of active users.[18] The former helps to emphasize the extremity of some of the spikes in number of tweets, and the latter shows the impact of spikes in users. For instance, in the normalized number of users per day graph, when the line jumps above 1.0, that means greater than 100% of active users are tweeting, i.e. twitter users who were previously inactive chose to log in again to tweet on that day. This can be used with the spikes in tweets data to see what days were particularly eventful for these accounts.

I ran a comparison of means test using these two samples for the mean of each variable count in each sample. With the way I ran the tests, a negative coefficient would indicate that there is a larger occurrence of the keyword in the three months leading up to my chosen event, and a p-value less than or equal to 0.05 indicates there is a statistically significant difference. Through this analysis, it was found that there is a significant increase in the usage of six of the 12 keywords: border, freedom, immigrant, maga, right, trump. This means that, for the sample of far-right Twitter accounts that I had available, the frequency of use of these keywords was noticeably higher before the white nationalist rally in Charlottesville. It also shows, however, that for these accounts, there was a significantly lower frequency of use of the terms 'nazi' and 'white' during the three months before the rally than during the period between March 2017 and

[17] Samantha Lachman, "White Nationalists Gather On Halloween To Discuss How Oppressed They Are," *Huffington Post*, Oct. 31 2015, accessed May 10, 2018, https://www.huffingtonpost.com/entry/white-nationalists-conference_us_56353a69e4b063179912ab4b.

[18] Reference Graphs 5 and 2.

January 2018, and that there was not a significant difference in the usage of the words Christian, genocide, holocaust, and pride.

This could be taken to indicate that there is in fact a correlation between the frequency of use of border, freedom, immigrant, maga, right, and trump and the occurrence of a far-right ideology driven violent event. The reverse frequency of the use of ‘nazi’ and ‘white’ could be explained by the public response to the Charlottesville event. The death of a counter-protestor at a white nationalist rally could have caused a backlash against white nationalist ideology from the general public, and therefore the accounts I gathered data from could have used the term less in order to avoid online persecution.

Moving Forward

One thing that could be done in the future is to collect and handle Twitter data in real time. While the Twitter API only allows 3200 tweets from any one account’s history, there is no limit that I have found to the number of accounts that can be tracked in real time, or the number of tweets that can be collected. With this, there is the potential to creating a script that would perform the variable counts discussed above on keywords. This could potentially be a way to predict a terror-type event rather than show a pattern before one has occurred in the past.

Another possible direction, and something I originally intended to do but did not have the time or skill, is to run sentiment analysis on the content of the tweets collected, again leading up to a terror-type event that has already occurred. Along this same thread would be to use the terror data collected and compare it against the different frequencies of keywords or the sentiment analysis and see if there is a correlation between the two.

Table 1	
Alt-right twitter users	FFloodgates, mikepolignano, RichardBSpencer, SnekNoExist
Anti-immigrant users	CarolineGlick, FAIRImmigration, glennspencer8, horowitz39, MarkSKrikorian, NumbersUSA, palehose67, rgarciaq, wwwCISorg
White nationalist users	AltRight_Voice, AngloFlaxen, charlesmurray, DailyShoah, davidlane, DrDavidDuke, EuropeanDefense, EvanMcLaren, GhostofJaredTay, Identity_SC, IdentityEvropa, JamesEdwardsTPC, JamieKelso, johndenugentESA, KyleBristow, Michael122448, NewRightAmerica, peterbrimelow, populistearl, Race__Realist, ramzpaul, RealHalTurner, realJohn_Wesley, rlhonoraryaryan, take_backUSA, TheMadDimension, ThreeGunsXd, TOOEdit, Y_A_Freedom
Keywords	white, trump, right, pride, nazi, maga, immigrant, holocaust, genocide, freedom, border, Christian

Table 2				
Keyword	Coefficient	Standard error	T	P-value
Border	-0.0086	0.001	-9.052	<0.0005
Christian	-0.0008	0.001	-1.460	0.144
Freedom	-0.0024	0.001	-3.320	0.001
Genocide	-0.0007	<0.0005	-1.648	0.099
Holocaust	-0.00002757	<0.0005	-0.129	0.897
Immigrant	-0.0047	0.001	-6.445	<0.0005
Maga	-0.0013	0.001	-2.007	0.045
Nazi	0.0018	0.001	3.553	<0.0005
Pride	0.00007373	<0.0005	0.521	0.603
Right	-0.0096	0.002	-5.963	<0.0005
Trump	-0.0387	0.002	-16.643	<0.0005
White	0.0037	0.002	2.369	0.018

Graph 1: Users per day

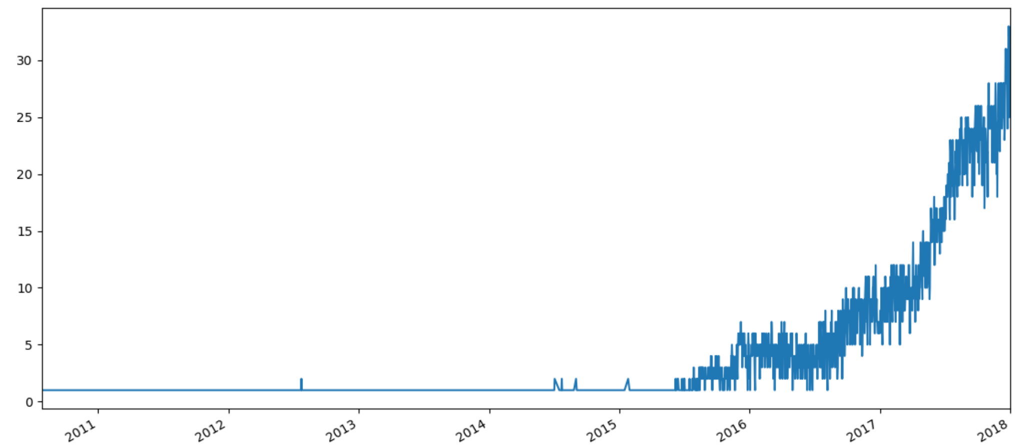

Graph 2: Normalized users per day

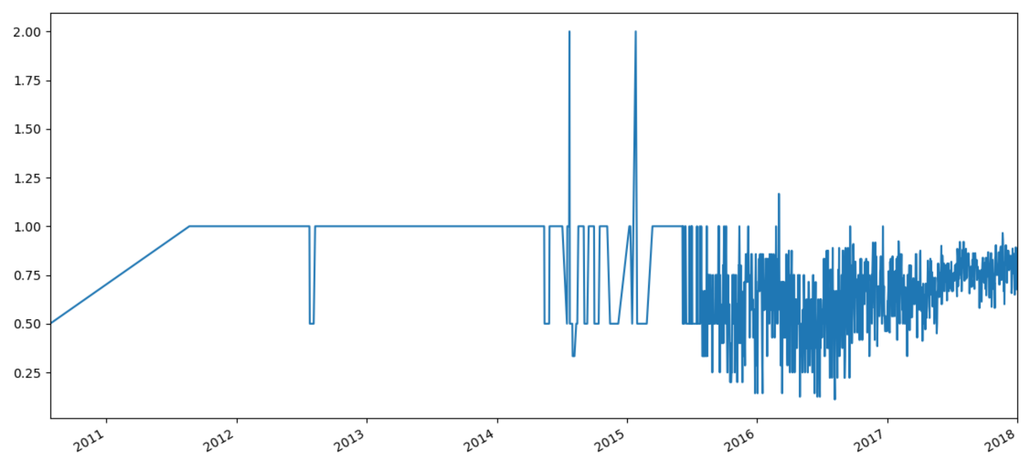

Graph 3: Active users per day

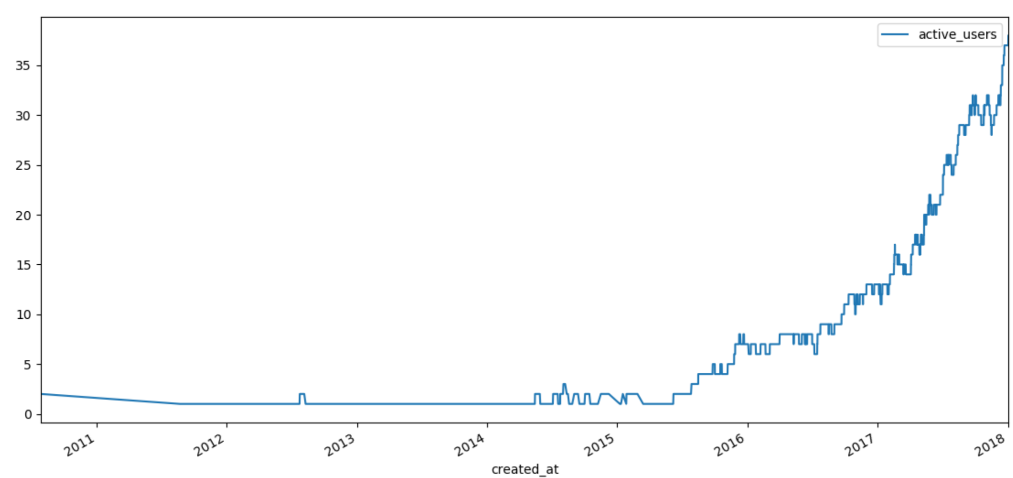

Graph 4: Tweets per day

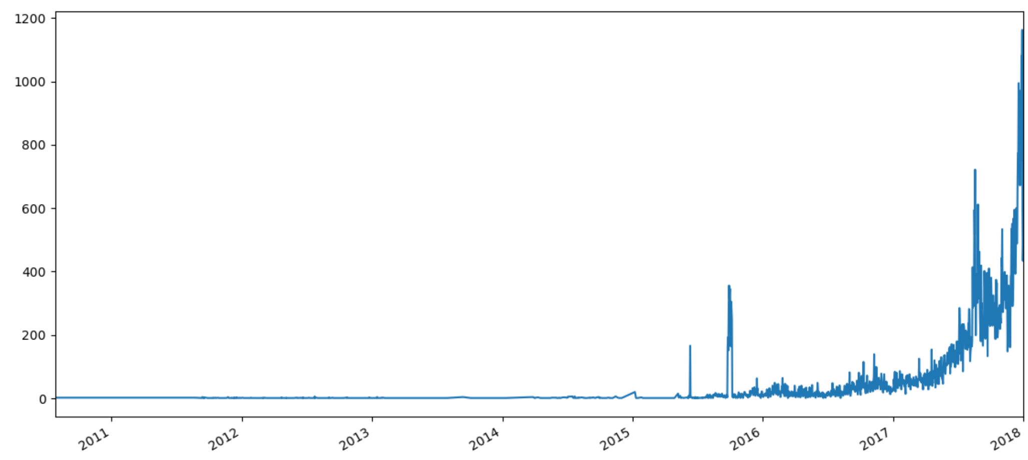

Graph 5: Normalized tweets per day

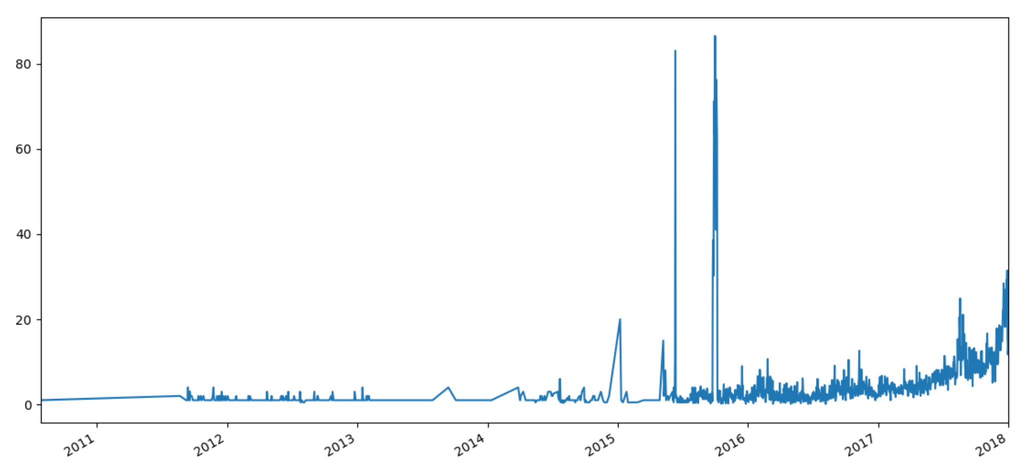

Bibliography

Brynielsson, Joel, Andreas Horndahl, Fredrik Johansson, Lisa Kaati, Christian Mårtenson, and Pontus Svenson. "Harvesting and analysis of weak signals for detecting lone wolf terrorists." *Security Informatics* 2, no. 1 (undefined): 11. http://dx.doi.org/10.1186/2190-8532-2-11.

Cohen, Katie, Fredrik Johansson, Lisa Kaati, and Jonas Clausen Mork. "Detecting Linguistic Markers for Radical Violence in Social Media." *Terrorism and Political Violence* 26, no. 1 (undefined): 246-56. http://dx.doi.org/10.1080/09546553.2014.849948.

Roof, Dylann. "Manifesto." *Last Rhodesian* (blog), June 20, 2015. Accessed May 10, 2018. https://www.documentcloud.org/documents/2108059-lastrhodesian-manifesto.html.

SPLC Southern Poverty Law Center. "Extremist Files." Accessed May 10, 2018. https://www.splcenter.org/fighting-hate/extremist-files.

Twitter Developer. "API Reference." 2018. Accessed May 10, 2018. https://api.twitter.com/1.1/statuses/user_timeline.json.

Yanofsky, David. "tweet_dumper.py." GitHub repository. Accessed January 2018. https://gist.github.com/yanofsky/5436496.