Using computational methods to analyze gang-related social media data

Megan Goodling
Columbia University in the City of New York
Natural Language Processing Lab
megoodling@davidson.edu

Abstract— Chicago has experienced a significant increase in gang violence, and this is in part due to social media sites such as Twitter, which serve as a means of communication and taunting among gang members. We collect and analyze this online data with the goal of creating tools for detecting when aggression manifests from the internet to the real world. Here I present an attempt to estimate the locations of gang-affiliated Twitter users for when geographical data is not available. I also explain my methodology in creating an Amazon Mechanical Turk program for extracting personal user data from Twitter profiles, while simultaneously educating youth annotators about the vulnerability of information they post online.

1 INTRODUCTION

Violence in Chicago severely spiked in 2016, as homicides rose 57% from the previous year. Gang related activity contributes to these numbers, and increasingly, gang members turn to social media sites such as Twitter for communication and expression. According to the Chicago Police Department, gang members use the platform to both incite violence with other gangs and mourn the loss of loved ones (“Chicago”).

This project focuses on the locations of gang-affiliated users. Since many users do not enable Twitter’s geolocation feature, concrete data about location must be manually extracted and interpreted. Subtle mentions of location are often found in tweet content, user handles, and user bios.

We investigate tweets from a gang networks across Chicago, which make up a vast dataset that provides insight into how gang members use social media. Some Twitter users post pictures of firearms captioned with threats, while others post about sadness and loss after a friend or relative is hurt or killed by a rival gang. Many engage in conversation about their feelings, other users, song lyrics, drugs, sex, and their whereabouts. The vernacular of the tweets in the dataset is quite far from standard English, so standard methods for analyzing text data must be adapted and generated accordingly.

2 DATA

The dataset used for this project is made of publicly available tweets from the network surrounding a formerly active Twitter user: Gakirah Barnes, who was a known Chicago gang member from age 13, until she was killed in 2012 at age 17. She was known as an assassin for her gang and posted over 27,000 tweets in a three year period (Blevins et. al, 2016). The tweets in the dataset come from her communicators and followers, and make up a set of about 2 million. They have been scored as conveying either aggression or loss by Blevins' program.

2 RELATED WORK

Desmond Patton and the Columbia University SAFE Lab have looked into the relationship between social media postings and violence, citing the 'digital street' as a growing space for youth expression that can occasionally manifest into action (Patton et. al, 2016). This research has been expanded upon by many others in the academic community.

Balasuriya et. al developed a method for detecting whether Twitter users are affiliated with a gang based on the vernacular with which they tweet. While gang members make up a small percentage of all Twitter users, the language they use is very distinct. Balasuriya uses a machine learning approach to detect these users amongst millions of other users. The program looked at features such as number of curse words and use of certain emojis to classify users, and the group experienced results of relatively high accuracy (Balasuriya et. al, 2016). However, the program does not determine what gang the users are affiliated with, which would be useful for analyzing gang dynamics and networks.

Additionally, Terra Blevins of Columbia University developed a part-of-speech tagger and phrase table, as well as a tweet classifier for aggression and loss in this specific vernacular (Belvins et. al, 2016). The tweets in dataset used in this project were automatically labeled by this classifier.

3 METHODS I

Since most users do not tweet with their geolocation on, the method we used for locating the users involved parsing through tweet content and user handles, looking for direct mentions of where they were. SAFE lab experts familiar with the vernacular directed us to look for this information in the form of blocks and cross streets; for example, #070 would represent 70th street. Alongside the search for location, we looked for any mention of gang names as well, with the hope that this would help us with grouping users. Again, members of the SAFE lab advised on which words referenced specific gangs, and these words were searched for in the dataset. Many Twitter users tweet about location and gang affiliation in various ways: some users' handles contain this information, and others mention this information in general tweets or conversations with other users. Each of these types of data were detected in the corpus and analyzed.

The hypothesis for the experiment was that a general sense of user locations could be determined based on the places they were talking about, and hopefully we would be able to visualize this in some helpful way. We hoped there would be enough data in the users' tweets and handles to accomplish this task. We also intended to look at individual users and clusters of users based on their mention of specific gangs.

The layout of the city of Chicago is very grid-like in nature. I took advantage of this fact in my primitive visualization of the data by plotting the relevant area with the python package matplotlib. The area was chosen based on guidance given by William Frey, a Columbia SAFE lab colleague; it spanned over the South side of Chicago from 31st to 91st street from north to south, and from S. Western Ave to Woodlawn Ave from west to east. A previously drawn map of gang territory was used a cross reference, as shown in Figure 1; the map was manually created by an anonymous social worker in 2015 and posted online. The color blocks represent the different small gangs across the South side. A possible goal of the automatic detection of tweets was to compare with and verify this map.

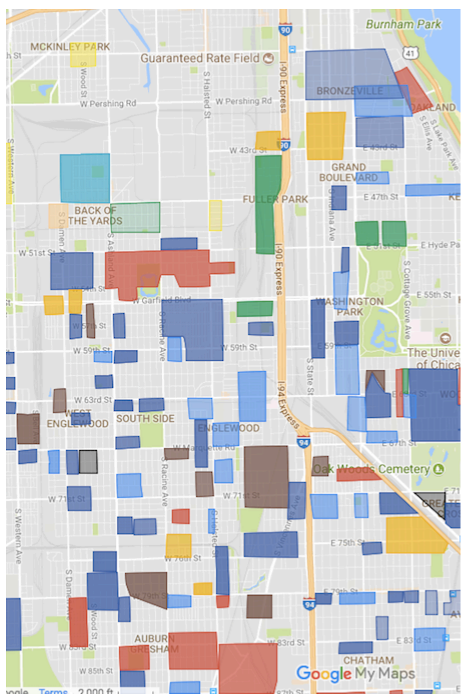

Figure 1: Manually drawn map of Southside Chicago gang territory

The aggression and loss scores generated by Blevins' was also factored in to my early visual representation. I represented aggressive tweets as red and loss tweets as blue.

In my program, each user was associated with all the streets, blocks, and gangs mentioned. A line was drawn on the grid at each detection of street name or block number. Each line was drawn very transparently so that darker lines represented a higher number of mentions. The original thought was that a grouping of many intersections of the lines might mean that the gangs were likely to occupy that general region.

4 RESULTS

When the program was run over all aggressive tweets from the set of 2 million tweets, the lines seemed to converge in the Southwest corner of the map. The area where the intersections of lines are most dense is near where the gangs were known to be located. However, it was not a very clear or definitive result, shown in Figure 2.

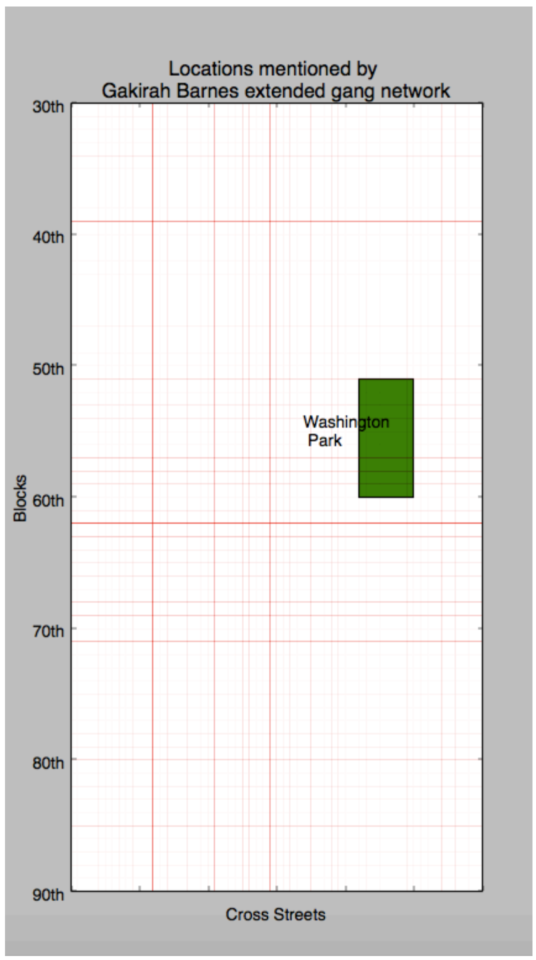

Figure 2: Mentions of locations detected in all aggressive tweets of dataset

The program was also run with the loss tweets (blue) and every single tweet (all green), with similar results.

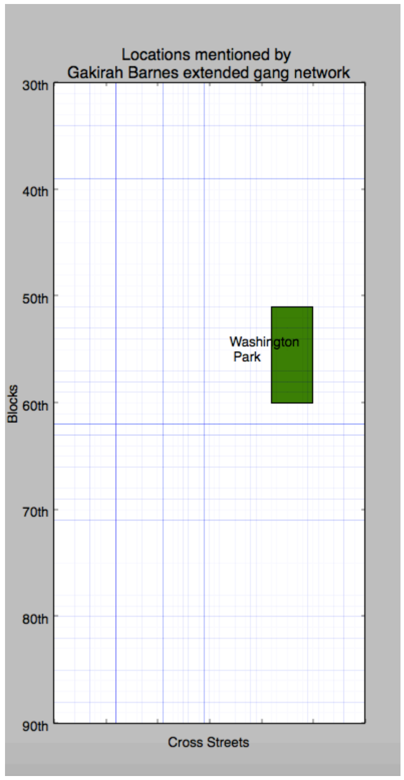

Figure 3: Mentions of locations detected in all loss tweets of dataset

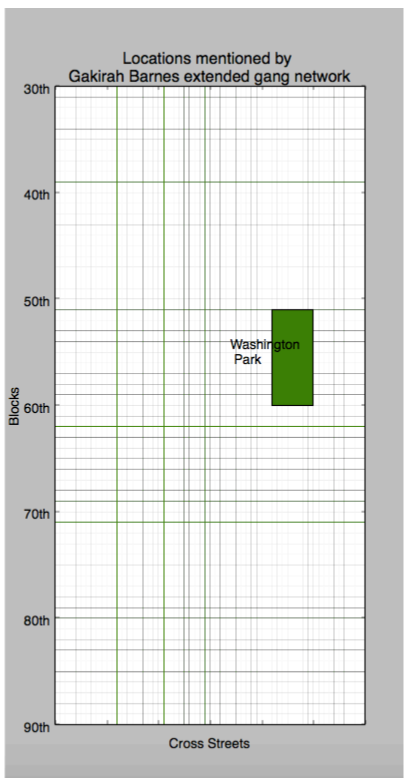

Figure 4: Mentions of locations detected in all tweets of dataset

When narrowed down to a single user, the plots of users that mentioned locations were often just two or three lines intersecting on the map, as shown in Figure 5. The purple lines represent a combination of blue and red overlaying lines, showing both aggression and loss from a single user.

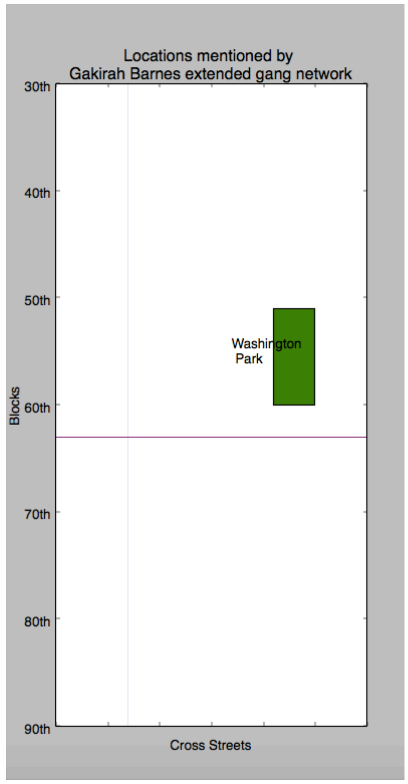

Figure 5: Locations mentioned by user @DOT_063

All of the visualizations generated looked similar to one of these patterns: either a series of many lines, or just a few lines with one or two intersections.

4 ERRORS WITH THE APPROACH

A major problem with this approach to the data was found when narrowing down the users to smaller subsets. Some users would only mention two or three locations. In some cases, these plots would just be a few horizontal lines spanning over many miles worth of the Chicago grid; this would not be helpful for determining an individual's location.

However, even if there was enough data to show several intersections, this would not necessarily represent the area being occupied. For instance, if a user was located at two specific locations and mentioned those street names, the cluster of intersections could be far from the user's actual locations.

A greater amount of data could potentially be helpful in improving accuracy of this fact, but further analysis and verification would need to take place.

Another significant problem with the code was that it referenced a dictionary of street names when searching for location. Sometimes, the street names were also commonly used words, such as "State" in State Street. Additionally, the block numbers searched for were often used in different contexts; for example, the number 69 is often used both in a sexual manner and as a reference to 69th Street. Thus, it was difficult to tell when users were talking about the physical location, and this greatly contributed to error by drawing lines over these areas when they did not represent a place.

A less significant error was the assumption that Chicago was a perfect grid; some streets do not abide exactly parallel to the others. However, this affected the data less than the other problems because it is for the most part grid-like and regular.

Furthermore, the program was limited in how it determined which gang to associate a user with. Users were simply associated with the gang they mentioned the most or the gang that is in their handle. Context was not taken into account, so if a user was volatile towards one gang and mentioned its name often, and did not mention his or her own gang, the user was associated with the rival gang.

5 METHODS II

Due to the abundance of errors with this approach, the group concluded that manually extracting mentions of locations from Twitter content and plotting on a grid is not an accurate way to determine gang member location. Thus, the group turned to a machine learning approach to the data to see if better analyses could be made.

We decided to create a data annotating project in which users would search for location and gang affiliation of users, as well as a host of other types of personal information, by looking at the profiles of Twitter users. In this initial phase, we aimed to create a pilot program in which four known users try out the annotation form and give feedback about how the process went. Another goal of the pilot program for annotating tweets was to implicitly educate our annotators about the large amounts of personal information that people unintentionally post to the web.

The data used for this experiment was a set of 200 links to profiles of users associated with Gakirah Barnes. These profiles were collected alongside the 2 million tweets used in the first part of the project.

My mentor suggested I use the Amazon Mechanical Turk, which I learned is an online platform for laypeople to annotate data posted by requestors. From the large amounts of data collected, machine learning algorithms can be developed. I designed a Human Intelligence Task (HIT) form that would allow annotators to click on a link to a user and fill in all personal data found, shown in Figure 6.

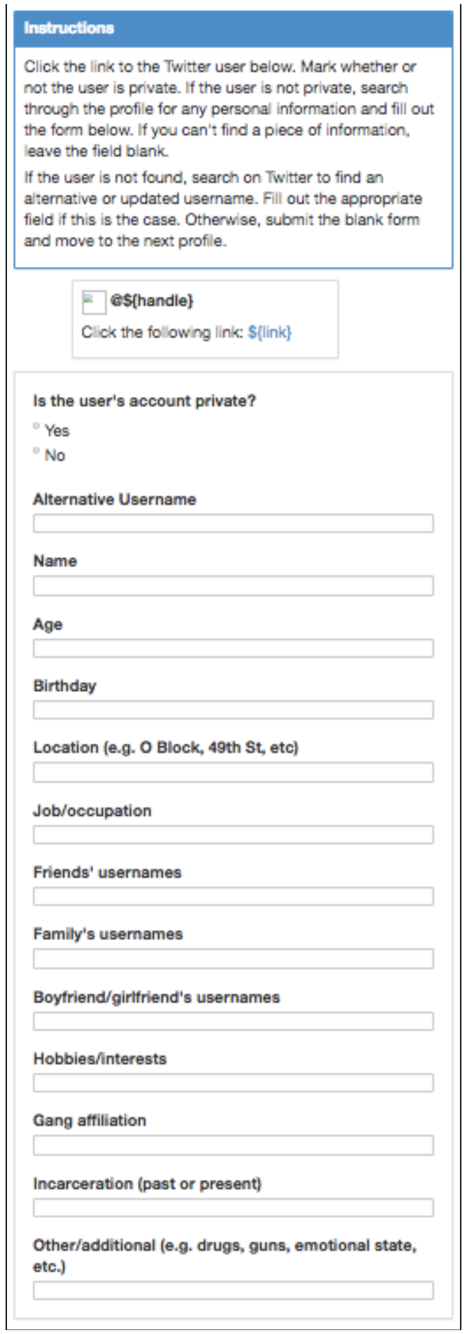

Instructions

Click the link to the Twitter user below. Mark whether or not the user is private. If the user is not private, search through the profile for any personal information and fill out the form below. If you can't find a piece of information, leave the field blank.

If the user is not found, search on Twitter to find an alternative or updated username. Fill out the appropriate field if this is the case. Otherwise, submit the blank form and move to the next profile.

@${handle}

Click the following link: ${link}

Is the user's account private?

Yes

No

Alternative Username

Name

Age

Birthday

Location (e.g. O Block, 49th St, etc)

Job/occupation

Friends' usernames

Family's usernames

Boyfriend/girlfriend's usernames

Hobbies/interests

Gang affiliation

Incarceration (past or present)

Other/additional (e.g. drugs, guns, emotional state, etc.)

Figure 6: HIT form for annotating a Twitter profile for personal information

We planned to present these HITs to a group of youth Summer Fellows working in the Digital Scholars Lab. I asked each of the Fellows to request approval to be an Amazon worker. Unfortunately, the Amazon worker approval process happened to be extremely backed up during the time of research. Thus, the workers were placed on a waitlist and were rendered unable to complete the data annotating during the time I was doing my research. However, this process taught me quite a bit about setting up data annotation. Alongside this, I researched and learned a lot about the actually machine learning algorithms and principles that take place once the data is labeled and collected. In the future I look forward to actually working with labeled data and exploring deeper into the field of machine learning.

6 CONCLUSION

From this project, my group concluded that accurate location of gang-affiliated Twitter users could not be accurately determined by parsing for mentions of location. More context is needed to accurately determine when users are talking about a location and where that location might be. While my mapping project did not yield any concrete results about the data, the group now knows that this approach will not suffice. The process was a great first exposure to research in that it allowed me to explore and manipulate data to see what types of visualizations and insights I could find. I experienced the research environment firsthand, attending research group meetings and exploring real-world data firsthand, and I learned a lot about current work being done in NLP.

I enjoyed becoming familiar with the Amazon Mechanical Turk platform and creating my first annotating project. This was my first exposure to machine learning, data annotation and mechanical turking, and I learned so much about each of these relevant computer science tools. I look forward to actually developing a machine learning program based on my results from another turking project one day.

ACKNOWLEDGEMENTS

I am grateful for all the support I received from my mentor, Dr. Kathleen McKeown, over the course of the project. I am also thankful for the DREU, Columbia University and my family for making this summer research experience possible.

REFERENCES

Lakshika Balasuriya, Sanjaya Wijeratne, Derek Doran, Amit Sheth. 2016. Finding Street Gang Members on Twitter. In ***WHERE PUBLISHED***. Kno.E.Sis, Wright State University.

Terra Blevins, Robert Kwiatkowski, Jamie Macbeth, Kathleen McKeown, Desmond Patton, Owen Rambow. 2016. Automatically Processing Tweets from Gang-Involved Youth: Towards Detecting Loss and Aggression. In ***WHERE PUBLISHED***.

Desmond U Patton, Jeffrey Lane, Patrick Leonard, Jamie Macbeth, Jocelyn R. Smith-Lee. 2016. Gang violence on the digital street: Case study of a South Side Chicago gang member’s Twitter communication. In *New Media and Society*, pages 2-3.

"Chicago Records 762 Homicides in 2016, up 57 Percent from Previous Year." *Fox News*. FOX News Network, 1 Jan. 2017. Web.