Title
Social Media and Crime Perception

Author
Feliciano, Matthew

Publication Date
2020

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Santa Barbara

Social Media and Crime Perception

A Thesis submitted in partial satisfaction of the

requirements for the degree Master of Arts

in Geography

by

Matthew W. Feliciano

Committee in charge:

Professor Keith Clarke, Chair

Professor Krzysztof Janowicz

Professor Somayeh Dodge

December 2020

The thesis of Matthew Feliciano is approved.

Somayeh Dodge

Krzysztof Janowicz

Keith Clarke, Committee Chair

December 2020

Social Media and Crime Perception

ACKNOWLEDGEMENTS

Foremost, my advisors have been invaluable in building out this thesis. Further, the DOLLY Project and the Floating Sheep Team at the University of Kentucky including Ate Poorthius, Matthew Zook, and Chen Qingqing were key by supplying the Twitter data and correspondence necessary to complete this research. Finally, yet importantly, I am always thankful to my family and friends for their support.

ABSTRACT

Social Media and Crime Perception

by

Matthew Feliciano

Expression of opinions are communicated on social media platforms, and these records can be analyzed and applied to a variety of spatial research questions. Using geo-tagged social media posts, specifically from the microblogging application Twitter, this research investigates the following research question: Assuming that users' perceptions of crime can be accurately extracted and analyzed from social media data, to what degree do they match with actual geo-tagged crime incidents in time and space? This research could be serviceable to officers of the law for effective administration of citywide resources with the mission to weaken crime in areas most affecting residents' perception of their safety.

The City and County of San Francisco was chosen for its availability of crime data and Twitter data. Methods applied to explore the posited research questions included preprocessing the Twitter content to remove unnecessary information such as URLs plus retweets. Sentiment analysis in RStudio separated the entire corpus of tweets into emotional categories. Finally, a series of point density and emerging hotspots maps were created to explore the relationships within the data. Comparison between the hotspots maps for the rates of crime and for the rates of tweets reveal different similarities and discrepancies based upon the time range used for crime. Based on the comparison between spatial and temporal patterns of crime and tweets showing fear, the interpretation is that the northeastern areas in

San Francisco were more likely to be the site of crime on Fridays around 6 p.m., plausibly larceny. Social media like Twitter may prove as effective indicators or perhaps even predictors of crime in space and time.

TABLE OF CONTENTS

LIST OF FIGURES

I. Problem Statement

Social media analysis is broad in scope. Platform options have ranged from Yelp, Facebook, Instagram, to Twitter, to name a few. This research focuses on the analysis of Twitter data for the specific purpose of analyzing the underlying emotion connotations locked within the tweets. This type of study is commonly referred to as *sentiment analysis* and is used for categorizing text by emotion (Liu, 2012). The current research addresses the question: is fear on social media related spatially and temporally to crime in similar areas? The hypothesis is that, in space and time, fear on Twitter would be highest in areas with the greatest reported crime.

There is a lack of research into the real-world events that affect specific emotions on social media. Many studies generalize sentiment into a polarity of negative or positive (Hatzivassiloglou & Wiebe, 2000; Turney 2001). Analyses of social media data were late on the timeline of geographic research (Hatzivassiloglou & McKeown, 1997) in part because studies that use the platforms as a source of data have only recently been proliferated by the internet and the subsequent boom of blog and microblog websites. Also, the usage of social media platforms in the United States have proliferated since 2005 when 5% of adults used at least one social media site while in 2019 72% of adults used at least one form of social media platform (Pew Research Center, 2019). The social media platforms eventually became mobile with the popularity of the smartphone and the dynamic applications that followed with its proliferation.

This research eschews general sentiment analysis for a more refined view focused on only one sentiment in particular, fear. Exploring new applications for methods is implicitly important research, however, this study could serve as a stepping stone to a future analysis for a deeper dive into the extraneous variables affecting the relationship between fear on social media and crime.

To test the hypothesis, data from a major social media platform, Twitter, is used as a source of text. The text is used for sentiment analysis via RStudio and the application of emotion classifiers accessible from the programming language's data packages and functions. The outcomes of this study are displayed in figures for crimes and tweets over time and are shown visually on maps to reveal spatially dense areas of both crimes and tweets.

The results of the sentiment analysis of tweets and the assessment of its patterns with crime could influence the allocation of resources to reduce fear in highly affected areas. Crime is documented by law enforcement so by extension law enforcement may be interested in the effects of crimes on the emotions of citizens.

II. Literature Review

A. Breadth of Applications

Social media platforms can be viewed as a source of big data for a range of applications, some of which are reviewed in this chapter. In particular, Twitter allows users to post 280-character messages in real-time called tweets and serves 152 million daily active users worldwide as of 2019, posting 500 million tweets every day (Clement, 2020; Smith, 2020).

In the political realm, social media has been used to gauge sentiment related to political parties and debates. One such study, Diakopoulos & Shamma (2010), built methods aimed at understanding the change in sentiment over time before, during, and after viewing political debates, particularly the first U.S. presidential debate in 2008. Using approximately 3,000 tweets linked by similar hashtags relating to the debate, the sentiment of the tweets was categorized as negative, positive, mixed, or other. These aggregated sentiments were used to determine which candidates the users advocated. The study showed that the general sentiment was negative and Obama rather than McCain was favored. Their study demonstrated that the overall sentiment of a debate can be categorized and the most supported candidate can be identified via social media platforms such as Twitter.

In behavioral economics, researchers measured the correlation between Twitter sentiment and the Dow Jones Industrial Average (DJIA) (Bollen, Mao, & Zeng, 2011). The study conducted sentiment analysis of ten million tweets using two sentiment analysis tools: OpinionFinder to measure the mood of the text as negative or positive, and Google-Profile of Mood States (GPOMS) to measure the mood in six terms: calm, alert, sure, vital, kind, and happy. The accuracy of the DJIA prediction models was improved with a high level of confidence using public sentiment measurements and a Self-Organizing Fuzzy Neural Network model. The study's accuracy in predicting the daily peaks and troughs in the DJIA closing values was 87.6%, which showed that dynamic public mood can be followed and that dynamic sentiment is responsive to social catalysts.

In the field of medicine, Yates & Goharian (2013) used 2,500 online consumer reviews to automatically extract adverse drug reactions (ADRs) that were not reported by the United States Food and Drug Administration (FDA). Their study generated a lexicon-based system able to identify patterns within the drug review forums. They pinpointed ADRs that were "expected" according to the FDA and pinpointed ADRs that were "unexpected", but commonly reported by consumers on social media websites dedicated to pharmaceutical drugs. The outcome of the study allowed for unexpected ADRs from breast cancer medications to be tracked, demonstrating the potential power of mining sentiment from social media sites.

B. Foundational Studies

In academia and within the agendas of policy makers, the attempt to measure the fear of crime has proliferated since the 1960s (Hale, 1996). Criminologists and the criminal justice system have been the main proponents for analyses of the fear of crime. According to Hale (1996), from the mid 1960's to the mid 1990's, over two hundred analyses including articles, conference papers, and books focused on the fear of crime. This means that theoretical and practical material on the subject has had ample time to be developed and honed.

Hale (1996) also examined the theoretical explanations and evidence for the causes behind the fear of crime by studying the physical, psychological, and economic vulnerability perceived by his subjects and the relationship to the level of crime (as a direct victim or indirect victim by word of mouth or media consumption). Two contributors were factored into the fear of crime: the personal perception of the risk of becoming a victim, and the assessment of how serious the repercussions will be to the person being targeted for a crime. "Incivilities become a symbol of decaying communities and may produce fear of crime even with relatively low objective crime levels" (Hale, 1996, p. 80). An individual's sentiment towards their neighborhood is a factor as well. Fear of crime carries real consequences such as people changing their habits to avoid victimization. Public places become off-limits or certain areas are avoided.

Box, Hale, & Andrews (1988, p. 341) concluded that "Because of its intrinsically disturbing nature and its adverse consequences for the quality of community life, fear of crime has become a major social problem." At the time that paper was written, surveys were the main form of collecting qualitative data, which is why social media is being considered by researchers as an alternative to garner large datasets at a lower cost. Examples of questions on these surveys include: How safe do you feel being out alone in your neighborhood after dark, or is there any place around here where you feel unsafe walking at night? But there are added implications to questions such as these, as pointed out by Ferraro & Grange (1987, p. 76): "A person who says he or she would not feel very safe may not be afraid at all, but simply aware of the relative risk." Thus, one may avoid walking alone in their neighborhood at night, but not manifest any fear of crime (Hale, 1996).

There has been significant advancement in the collection of new data types and subsequent analyses within the past decade. Pak & Paroubek (2016) performed a linguistic analysis using a sentiment classifier to determine positive, negative, and neutral sentiments for 300,000 tweets. Their study categorized the sentiments further, into happiness, amusement, or joy for positive emotions and sadness, anger, or disappointment for negative emotions and neutral for those tweets that only stated a fact or did not express any emotions.

Other studies such as Pang & Lee (2008), believed the progression of research into the realm of social media was a natural progression because understanding what other people think is a key component of human behavior. According to their surveys, approximately 60% of Americans have researched a product online, which influences approximately 75 – 90% of consumer purchases. Their conclusions were that the year 2001 marked the beginning of sentiment analysis and opinion mining, and the contributing factors included the rise of machine learning in natural language processing and the availability of datasets due to the proliferation of the internet.

More specific to the current research topic, a study by Chainey, Tompson, & Uhlig (2008) used hotspot mapping to predict spatial patterns of crime. According to the study, hotspot mapping is inherently a form of crime prediction as it relies on past data to identify high crime areas to which police and resources should be directed to reduce crime. Multiple methods for mapping hotspots are available such point mapping, thematic mapping, standard deviational ellipses, and kernel density estimation (KDE) (Chainey, Tompson, & Uhlig, 2008). Hotspots afford the additional capability to not only identify locations of interest, but also the size/shape and orientation (Chainey, Tompson, & Uhlig, 2008). Four types of crime were compared: burglary, street crime, theft from vehicles, and theft of vehicles. The conclusion was that kernel density estimation was the technique that outperformed standard deviational ellipses and grid-based thematic mapping. Through hotspot mapping with kernel density estimation, the study was able to identify hotspots and their proximity to structures such as shops and markets. Street crime remained fairly static in space so hotspots from past data could be good indicators of crime in the future. However, more transient crime like vehicle theft were less reliable when considered in the same fashion.

C. Recent Research

It is pragmatic to consider the schema of sentiment analysis. First, the goal of sentiment analysis is to identify the emotions behind the text of interest, and there are two major methods: machine learning or lexical (Gonçalves, Araújo, Benevenuto, & Cha, 2013). Machine learning is good at adaptive classification of sentiment for analyses specific to potentially niche study areas, however, it requires annotated data that is not commonly available in open source, forcing researchers to pay a premium for a more accurate machine learning analysis (Gonçalves, Araújo, Benevenuto, & Cha, 2013). Lexical sentiment analysis, in contrast, uses a list of words preassigned to a sentiment, making for a broader review of sentiment, but is more cost-effective (Gonçalves, Araújo, Benevenuto, & Cha, 2013).

Detecting the fear associated with crime is closely related to documented crime itself. Wang, Gerber, & Brown (2012) provided a preliminary report on the efficacy of using Twitter to predict crime incidents. They used automatic sentiment analysis and latent Dirichlet allocation to predict crime using a linear model. The aim of the study was to predict hit-and-runs that would occur in the future. One potential error that the study reported was the assumption that the events contained in the tweets were those that occurred on the same day as the posting.

Intravia, Wolff, Paez, & Gibbs (2017) looked specifically at the relationship between social media usage and the fear of crime. Their aim was to study the relationship between several forms of social media and the fear of crime in addition to a second objective to determine whether the relationships differed when the audience's characteristics were changed such as age, race, sex, and prior victimization. However, it is worth noting that the majority of social media users are younger males (Clement, 2019; Aslam, 2020), but the importance of the platform is reinforced by the statement that roughly 66% of American adults get part of their news from social media (Moon, 2017). A main hypothesis in Intravia, Wolff, Paez, & Gibbs (2017, p. 165) was that social media consumption is positively correlated with the fear of crime: "Overall social media consumption and fear of crime is conditioned by perceptions of safety." Their findings supported the idea that the ingestion of social media may have stronger effects for people without direct experience with crime, which can be stated differently: people who feel safe are more susceptible to being induced into fear by consumption of social media.

The following research will pick up from the steps taken in prior investigations, address some of the shortcomings mentioned, and target a city previously unstudied for sentiment analysis of social media and crime, San Francisco, California.

III. Data Collection and Processing

A. Data Considerations

The crime datasets and the crime categories used in this research were adherent to the Federal Bureau of Investigation's criteria per the UCR Handbook (FBI, 2004). The crime datasets were open source and easily accessible online through DataSF. This was not the case for the Twitter data set. The tweets were supplied by the Floating Sheep team working on the DOLLY Project and as such it was requested that it not be further disseminated. Per this agreement, the Twitter dataset was only to be used for the purposes of this research. The crime dataset and the Twitter dataset were crucial in working towards discerning any potential correlation between fear on social media and real-world crime incidents.

B. Data Sources

The data used in this study was acquired through the DOLLY Project (Digital OnLine Life and You). DOLLY is a data repository of historical geotagged tweets made possible by the Floating Sheep team. The DOLLY Project customized existing open source technology to create a back-end that collects roughly eight million tweets per day, completes foundational analysis, and finally indexes and geocodes the corpus to allow for real-time queries. Their project started collection in December of 2011. The corpus of tweets provided by the DOLLY Project for San Francisco totaled 217,457 geotagged tweets from June 21st, 2012 to March 27th, 2018.The dataset included the Twitter user identification number, the date and time that the tweet was created, the coordinates of the tweet, and the text of the tweet.

DataSF is the City and County of San Francisco's official open data portal and was used as the source of the documented crime incidents. DataSF launched in 2009, and holds hundreds of datasets for application by researchers and residents alike. The data acquired from DataSF included 2,522,433 geolocated crime incidents with the date, time, and crime category from January 1st, 2003 to the present (March 2nd, 2020 at the time of collection). Both the Twitter and the crime data sets were entirely in English.

C. Data Preparation

The tweets supplied by the DOLLY Project data repository were already indexed for date, time, text, and location in a comprehensive package over the span of several years. In contrast, there were originally two crime datasets, one from 2003 to 2018 and one from 2018 to present, which had variations in column content. In order to merge the two into one useable data frame, preliminary refinement was completed. Essentially, all columns in both datasets were given the same names where applicable and include: category of crime, police district, coordinates, and timestamp. Rows in which no data were available for the crime category were removed, incidents labeled “non-criminal” were removed, and categories were renamed for consistency (i.e. “battery by juvenile suspect” to simply “battery”).

D. Inherent Errors

In terms of spatial and temporal accuracy of the Twitter and crime data, there were some key differences. The geotagged tweets within the corpus collected from San Francisco yielded the exact longitude and latitude location as well as the date and time down to the second at which the user posted the tweet. While the tweets were automatically tagged in this manner when created, the documented crime incidents were not. It is posited that there was inherent human error involved with the creation of the crime dataset. The cataloging of the date and time is up to the officer to record and while the date most likely remains accurate, there cannot be exact time recordings because the incidents were not automatically recorded by a system, but rather by an officer, or reported by an eyewitness or victim after the fact. This implication may not be important for small scale analyses, but for analyses involving specific hours of the day, the potential error could have cascaded. Further, many of the locations for the crime incidents were recorded in terms of latitude and longitude of road intersections. This means that many of the incidents were most likely not spatially precise, but rather fell within a buffer zone of roughly the entire area of the road intersection.

The crime dataset raised some more obvious flags in terms of potential or implicit errors. In particular, the crime category totals hint towards incomplete data. For example, the dataset comprising 2003 to the end of 2017 is missing "homicide" incidents, which affected the distribution of that crime type from the previous decade. The same held true for other crime types such as "malicious mischief." Crime categories such as "other offenses" and "other miscellaneous" were combined into one "miscellaneous" category and removed because the incidents were too varied to hold much meaning in the analysis. Also, one might expect certain crime categories such as "drug offense" to decrease in California as marijuana was recreationally legalized in late 2016 and therefore would not be recorded as a crime incident as frequently since then. "Aggravated assault" and "battery" were combined into "aggravated assault" to fit the Part I crime definition as described by the FBI's Uniform Crime Reporting (UCR) Handbook.

Within the Twitter data, there were retweets, tweets with similar or nearly identical text, and there was the potential for tweets in the corpus to have been created by bots, law enforcement accounts, or entities/accounts used by multiple individuals. Some of these potential errors were remedied and will be reviewed in the following chapter.

IV. Methodology

A. Twitter

The initial step in the sentiment analysis of tweets targeted at fear was to collect a group of keywords to sift out unwanted tweets. To this end, Python's "Newspaper3k: Article Scraping and Curation" (https://newspaper.readthedocs.io/en/latest/) package was employed. The package is capable of scraping text within an article such as an online newspaper story, extracting information such as keywords, and summarizing the results. The Newspaper3k package was used to extract an initial set of 25 keywords from ten online Los Angeles Times articles on various forms of crime ranging from homicide to hate crime. A set of 8,784 tweets within the Los Angeles area containing these keywords was obtained using the web service Crimson Hexagon (previously called Brandwatch: https://www.brandwatch.com/). The keywords were pinpointed from several crime-related newspaper articles from the LA Times because of its wider audience than the San Francisco Chronicle (Agility PR Solutions, 2015). Following this, the Natural Language Toolkit (NLTK) (https://www.nltk.org/) suite of libraries was used to facilitate Latent Dirichlet Allocation (LDA) on the corpus of ten thousand tweets collected from Crimson Hexagon. LDA can be applied to create models that discern key topics across the entire corpus of text or individual text document (Tasci & Gungor, 2009). Using LDA, a bag of words was determined, which comprised the topics found over the entire corpus of tweets. This list of topics was the final set of keywords supplied to the DOLLY Project and Floating Sheep team. The keywords are as follows: abuse, assault, attack, child, crime, drug, force, gang, injury, kill, murder, rape, robbery, safe, sexual, shoot, suffer, suspect, theft, vandal, violence, violent, wound. Along with the list of keywords and the bounding box for San Francisco, the final corpus of tweets was compiled for sentiment analysis and correlation with the crime incidents.

The next step was to preprocess the tweets prior to sentiment analysis. The text within the tweets were refined in Python using the “Beautiful Soup 4” (bs4) (https://pypi.org/project/beautifulsoup4/) and NLTK packages. The bs4 package is capable of extracting data from XML files as well as parsing text into words and word stems. The NLTK and bs4 packages were used to remove any text elements that were non-alphanumeric, to remove any web links based upon the presence of “http”, and to remove any retweets based upon the usage of the “@” symbol. In addition, stopwords such as “I”, “me”, “a”, and “to” were removed based upon the corpus within the “stopwords” library of the NLTK package.

Following preprocessing, the tweets were imported into RStudio, an integrated development environment for R, for sentiment analysis. Prior to the sentiment analysis, twitter accounts within the corpus that contributed to 100 or more tweets were reviewed for verification of the type of user. This meant that accounts associated with job advertisements or news stations were removed because civilian users and their tweets were the target for the sentiment analysis. Once completed, the “syuzhet” (https://cran.r-project.org/web/packages/syuzhet/vignettes/syuzhet-vignette.html) package, which holds four sentiment dictionaries and methods to extract sentiment using these dictionaries, was employed. Specifically, Saif Mohammad’s “NRC” (https://saifmohammad.com/WebPages/NRC-Emotion-Lexicon.htm) emotion lexicon was used to extract sentiment, which is an unsupervised approach to emotion classification. The “NRC” emotion lexicon is an inventory of words and their relationships with eight emotions including joy, sadness, disgust, surprise, trust, anger, anticipation, and fear, all of which can be separated into two major sentiments, negative and positive.

Sentiment analysis for the entire corpus of tweets was conducted, and the emotion level for individual words in the Twitter corpus as well as for each tweet was assessed. However, the focus of further analysis was the emotion of fear for each tweet, which required selecting only the tweets that have the emotion of fear assigned to the majority of words in the text. Tweets with two or more emotions representing the majority of words in the text were automatically tie-broken in a statistically random fashion in RStudio.

The timestamp for each tweet was updated from Universal Time Coordinated (UTC) to the local San Francisco time (UTC-7). Plots for the tweets showing fear were created. For comparison, similar plots were constructed for the entire corpus of tweets, regardless of classified emotions. Map layers were produced from these datasets in the World Geodetic System 1984 datum (WGS 1984), which were later projected to the Teale Albers projection using WGS 1984 to work in distance measurement units, instead of geographic degrees.

Several maps were created for all the tweets acquired from the DOLLY Project and the tweets associated with fear, which included point density and emerging hotspot analysis. Point densities were used to show areas with the most recorded crimes and tweets in a spatial sense. Emerging hotspots were used to identify the temporal trends in the clustering of crimes and tweets; areas were categorized as a form of hotspot or coldspot within a space-time cube. The space-time cube was created by aggregating points into space-time bins based on the parameters inputted such as the point data, timestamp, time step interval, time step alignment, distance threshold, and aggregation shape type. For each bin, the Getis-Ord Gi* statistic calculates the z-scores and p-values for areas with high or low values of spatial clustering and functions by considering each point feature relative to neighboring features

(Getis & Ord, 1992). Upon completion of the calculations of z-scores and p-values, each bin is classified as a hotspot or a coldspot. Statistically significant positive z-scores are hotspots, and the larger the z-score, the greater the intensity of the hotspot (Getis & Ord, 1992). In contrast, statistically significant negative z-scores are coldspots and the smaller the z-score, the greater the intensity of coldspots (Getis & Ord, 1992). Then, the hotspots and the coldspots are analyzed in an independent time-series for each bin, the Mann-Kendall trend test, according to Esri's Hotspot Analysis (Getis-Ord Gi*) toolkit. The Mann-Kendall test is a rank correlation analysis that functions by comparing the bin value for the first time period to the bin value for the second, and if the first bin is smaller than, greater than, or equal to the second, the outcome is a +1, a -1, or a 0, respectively (Hamed, 2009). For each sequential pairing of time periods (one month to the next in the current research), the results are summed (Hamed, 2009). The expected sum is zero – no trend. To determine statistical significance for each sequential pairing of bin time periods, the difference between the expected sum is compared to the observed sum along with the variance for the values, the number of ties, and the number of time periods (Hamed, 2009). A z-score and a p-value is reported for the trend for each bin time series. Finally, each study area location (hexagons in the current research) is categorized as one of 16 types of hotspots or coldspots or as having no detectable pattern.

Multiple radii, distance thresholds, and time step alignments were tried as input for the point density and hotspots maps, however, the focal areas remained consistent throughout. Radii and distance threshold values of 50 meters to 1000 meters and time step alignments at the beginning and at the end of the time step intervals were tested for the point density and hotspot maps. The point density maps (https://pro.arcgis.com/en/pro-app/tool-reference/spatial-analyst/point-density.htm) for tweets and crimes were produced with a radius of 500 meters, roughly three to five city blocks. San Francisco is a dense city spanning a relatively limited area for the total population living there. Therfore, three to five city blocks allowed map visualization to be produced with a granularity that was generalized for the distribution of tweets throughout the city at a coarser scale. The emerging hotspot analyses (https://pro.arcgis.com/en/pro-app/tool-reference/space-time-pattern-mining/emerginghotspots.htm) for tweets and crimes were conducted using a space-time cube with a distance threshold of 250 meters and a time step interval of one month with hexagon grid bins and an "end time" time step alignment. The distance threshold of 250 meters was chosen to allow for a more fine-grained view of the hotspots, in contrast to the general view afforded by the 500-meter radius in the point density maps, and time step intervals of one month were chosen. . The hexagon bins were chosen because they are preferred in distributions that are affected by connectivity and/or movement paths such as roads (Birch, Oom, & Beecham, 2007), which are dense in the urban setting of San Francisco. An "end time" time step alignment was chosen because the focus was on the most recent events in the corpus, which allowed for a better view of the most recent emerging hotspots and placed any bias towards the beginning of the time frame.

The emerging hotspots were enhanced with additional attributes from the "Enrich" tool (https://pro.arcgis.com/en/pro-app/tool-reference/analysis/enrich.htm) in ArcGIS Pro, to include the 2010 total population and the total number of smartphones per individual. For each dataset, rows with no data for the 2010 total population were removed as those areas could not be used to calculate a population-based rate for the next part of this study.

Succeeding enrichment, tweets in the entire corpus acquired from the DOLLY Project and tweets showing fear were calculated as a rate based on the 2010 total population. This was completed in RStudio using the "EBest" function (https://www.rdocumentation.org/packages/spdep/versions/1.1-3/topics/EBest) in the "spdep" package (https://www.rdocumentation.org/packages/spdep/versions/0.1-2) to calculate the Empirical Bayes smoothed tweet rates. The "EBest" function in RStudio uses an empirical Bayesian estimation determined by the methods of moments, which are estimation methods similar to least squares, instrumental variables, or maximum likelihood (Chaussé, 2010). The methods of moments are the estimates of the statistical population parameters and are used for the smoothing of the rates, which reduces noise from the dataset (Pobuda, n.d.). Smoothing is conducted through calculation of the z-score, whereby the population mean is subtracted from each raw rate value and divided by the standard deviation (Pobuda, n.d.). The results from the Empirical Bayes smoothed tweet rates were then depicted in RStudio with 90 – 99% confidence hotspots to show the areas in the city with the greatest tweets normalized by total population.

B. Crime

The crime data obtained from DataSF were analyzed in RStudio. Plots of crime incidents and the types of crimes committed are provided in Chapter V. Plots were made for all crime in the dataset as well as specifically for Part I crimes, which are eight offenses that are serious and occur frequently, including criminal homicide, forcible rape, robbery, aggravated assault, burglary, larceny-theft, motor vehicle theft, and arson (FBI, 2004). As with the Twitter data, maps were produced from the crime datasets in the WGS 1984 datum, which were later converted to WGS 1984 Teale Albers to continue working in a similar datum, but transition to a projection that can be used for the radii in the point density maps and can be used for the distance thresholds in the hotspot maps.

In ArcGIS Pro, maps, displayed in Chapter V, were created for all crime incidents across the entire time range, for all crime incidents within the same time range as the tweets, and for Part I crime across the entire available time range and the same time range as the tweets. The same methods used for the Twitter data were applied for the crime data and similar products in term of plots and maps were produced.

The methodology for the crime data sets was not developed or explored extensively because the focus was not on the intrinsic characteristics in space and time of crime in San Francisco, but on the relationship crime incidents in the city have with tweets. In this sense, the analysis of crime was kept rudimentary and focused mainly on the total counts of crime across time throughout the city.

V. Results

A. Twitter

Total emotion scores for all tweets in the corpus (tweets based on keywords) are represented in Figure 1. The sentiment category with the greatest percentage was "fear" comprising 26.86% of tweets in the corpus – an unsurprising result considering the keywords provided to the Floating Sheep Team for acquisition of tweets from the DOLLY Project. However, the antitheses to fear, trust and joy, were still relatively high.

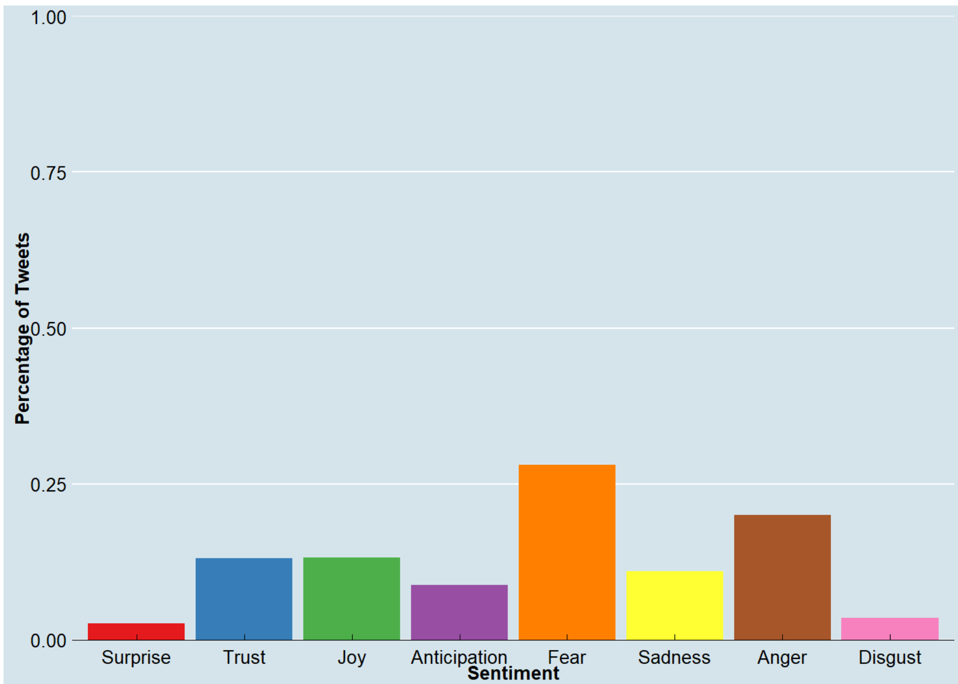

Fig. 1. Sentiment Scores for All Tweets – Emotion percentages for all tweets within the preprocessed corpus.

A view of all the tweets in the corpus (tweets based on keywords) and tweets showing fear is shown in Figure 2. There are two maxima in February (15,308 for all tweets and 4,361 for tweets showing fear) and June of 2016 (17,822 for all tweets and 5,235 for tweets showing fear) with a steep decline and incline between the two in April (8,799 for all tweets and 2,484 for tweets showing fear). Eschewing years with incomplete data, 2013 shows the lowest totals (18,764 for all tweets and 5,177 for tweets showing fear) compared to the year with the greatest total, 2016 (67,061 for all tweets and 18,987 for tweets showing fear). The consistent similarities in total tweets in the corpus and total tweets showing fear through time indicates that fear on Twitter is not a significant issue at this broad time scale in San Francisco.

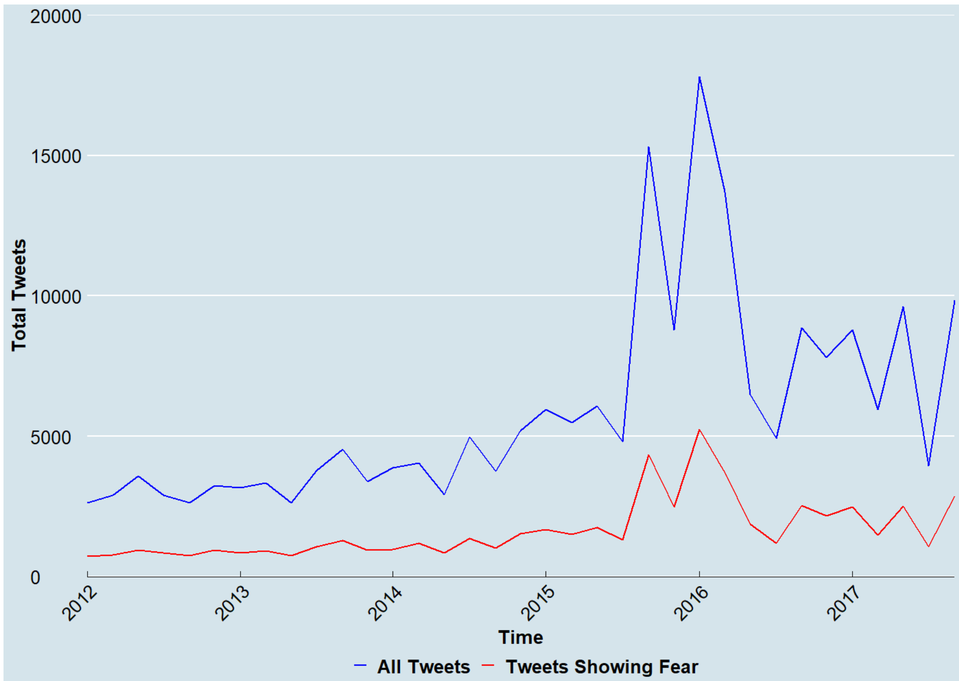

Fig. 2. Tweets Over Time – Totals for all tweets and tweets showing fear over the entire time range of the preprocessed corpus. Each year marker falls on June 1st, the first month represented in the corpus. Note 2018 is not marked because the last date in the corpus is March 27th, 2018.

Boxplots with the means for all tweets per day of the week and for tweets showing fear per day of the week show the same pattern (see Figure 3). Sunday has the lowest number of tweets (2,652 for all tweets and 959 for tweets showing fear) while Friday has the highest (3,501 for all tweets and 975 for tweets showing fear). There is not a significant difference in the patterns between the means for all tweets per day of the week and for tweets showing fear.

The overall tweets/tweets showing fear per hour of the day show some similarities (see Figure 4). For both, hour 10 is the absolute minimum on a given day (933 for all tweets and 266 for tweets showing fear), and the hours between 9 and 13 are an overall trough. For all

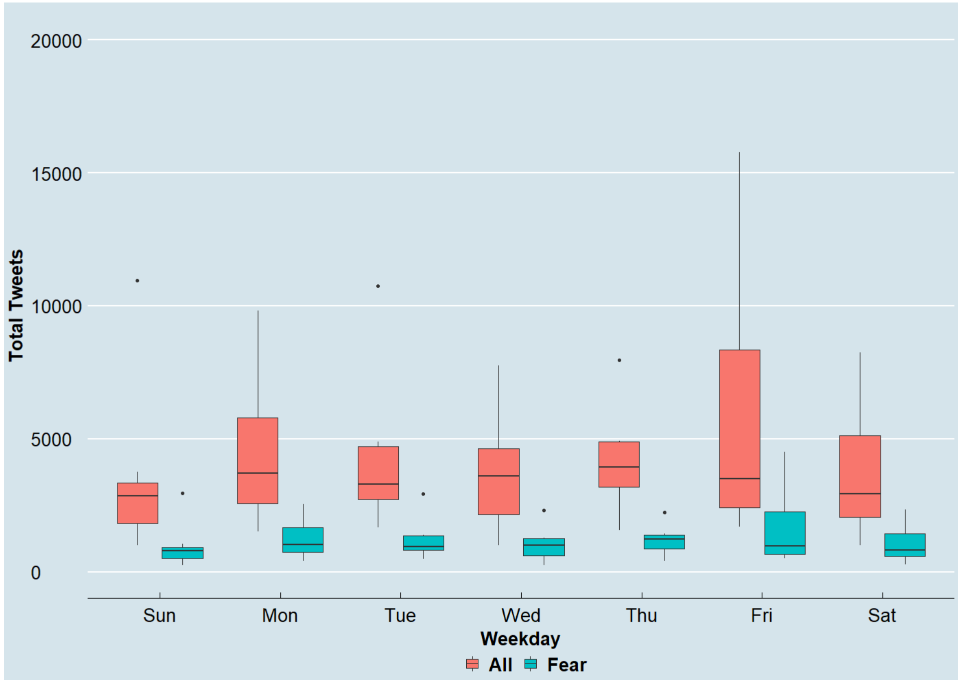

Fig. 3. Tweets Per Day of the Week – Means for all tweets and tweets showing fear per day of the week.

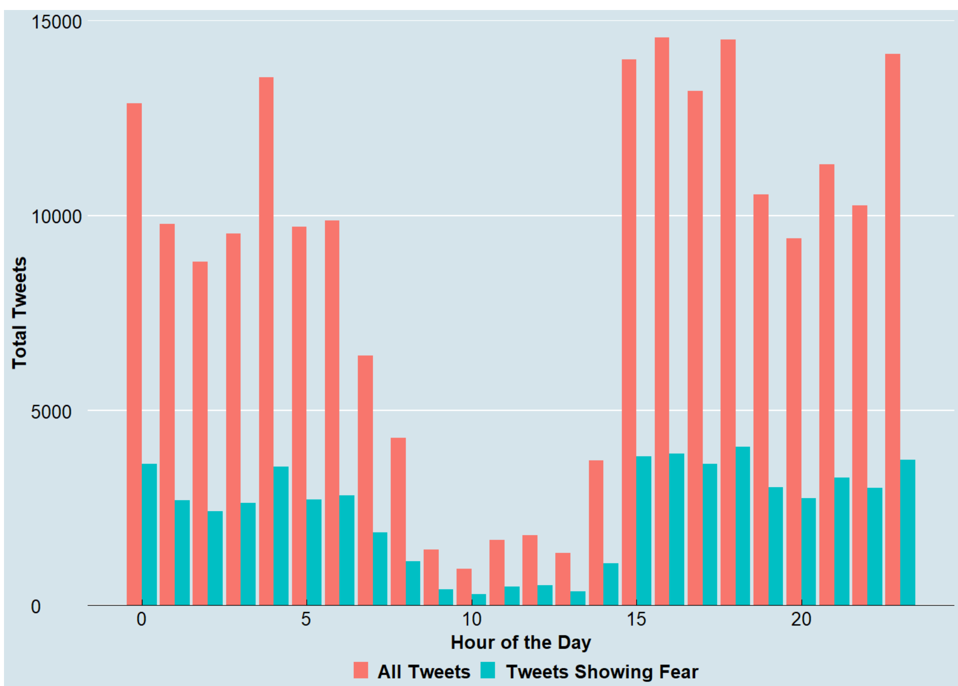

Fig. 4. Tweets Per Hour of the Day – Totals for all tweets and tweets showing fear per hour of the day.

tweets, the absolute maximum occurs at hour 16 (14,568 tweets) while in the tweets showing fear, the absolute maximum is hour 18 (4,063 tweets). There is no significant difference between overall tweets and tweets showing fear per hour of the day.

Point density can be useful for understanding the spatial distribution of all tweets and the tweets showing fear (see Figures 5 and 6) throughout the city of San Francisco. The spatial pattern of point densities between all tweets in the corpus and tweets showing fear remain similar. There is a large area of point densities focused on the Haight-Ashbury neighborhood, and the majority of the other relatively high-density areas are clustered around the northeastern sections of the city, mainly the Tenderloin, South of Market, and Financial District/South Beach neighborhoods. However, in the southwestern area of the city, there are greater point densities of crime in the Lakeshore neighborhood.

Point density maps were useful for a wide view of the crime snapshot in San Francisco, but hotspots and coldspots are more useful for recognizing the temporal variations. In light of this, the hotspots maps are attributed several categories for a deeper analysis into the progression of tweets showing fear in San Francisco. The categories and descriptions are important to keep in mind when viewing the hotspots maps. As previously mentioned in Chapter III, each time step interval spans one month and remains so in all subsequent emerging hotspots analyses for tweets or crimes Again, for all emerging hotspots analyses for tweets and crimes, the time range from the beginning to the end of the analysis is the full range available for the dataset. The last time step interval concludes at the end of March 2018 to align with the date range of the Twitter data.

The emerging hotspots for all tweets and tweets showing fear (not visually depicted in this research) are categorized mostly as sporadic with many on the western side of the clustering, and the rest of the areas are consecutive hotspots on the eastern side of the

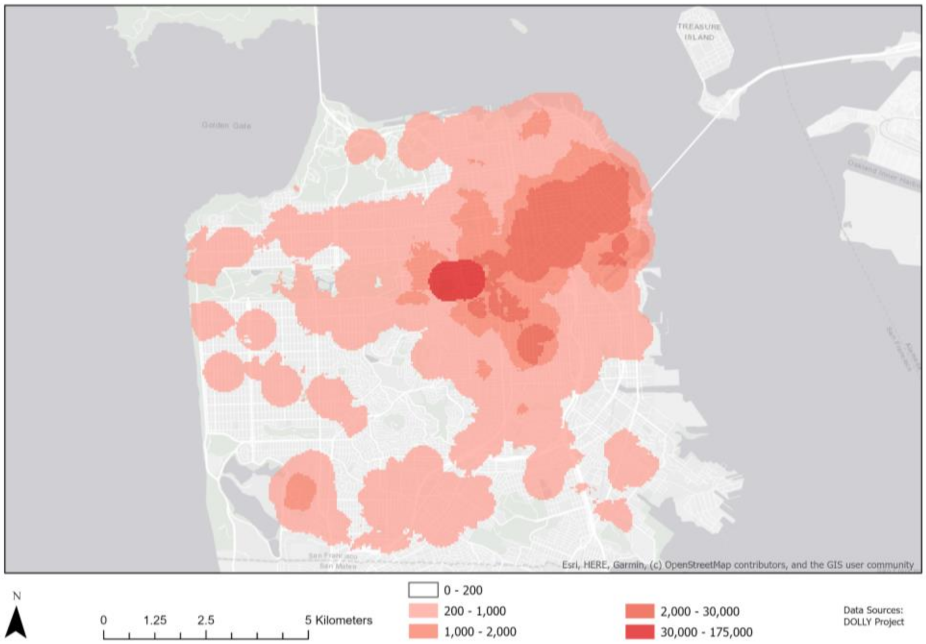

Fig. 5. All Tweets Point Density Map – San Francisco areas with low to relatively high concentrations of tweets. Note the artificial circles due to the radii distance interpolation method.

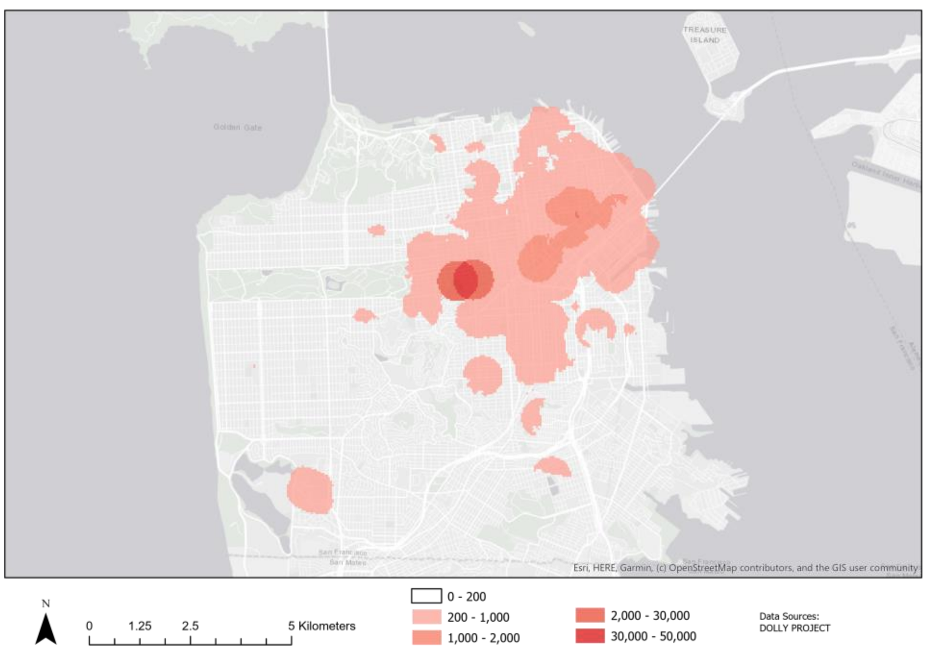

Fig. 6. Tweets Showing Fear Point Density Map – San Francisco areas with low to relatively high concentrations of tweets showing fear. Note the artificial circles due to the radii distance interpolation method.

clustering. The sporadic hotspots have been on-and-off hotspots over the entire span of the time range, but less than 90% of the time step intervals are statistically significant hotspots. Consecutive hotspots have a single unbroken sequence of statistically significant hotspots in the final time step intervals, but have never been statistically significant hotspots before. The hotspots for tweets showing fear were almost entirely consecutive hotspots with only two hexagons categorized as sporadic hotspots.

The hotspots per total population as a rate are shown in Figures 7 and 8. A relatively large cluster of 99% confidence hotspots center on the Haight-Ashbury neighborhood and another cluster of mostly 99% confidence hotspots span the South of Market and Financial District/South Beach neighborhoods. There is also a periphery of hotspots along the coast of the North Beach neighborhood. There are not significant differences between the hotspots for all the tweets in the corpus and for tweets showing fear.

Figure 9 shows smartphones per population as hotspots. Only 99% confidence hotspots were present, which span the Financial District/South Beach, South of Market, Mission Bay, and Potrero Hill neighborhoods along the far east side of the city. The point density maps and the hotspots maps for both all tweets in the corpus and tweets showing fear show areas of the city that overlap with the hotspots shown in Figure 9 – around the Mission Bay neighborhood. The Haight-Ashbury neighborhood was not significantly represented by hotspots of smartphones yet was shown to be a hotspot for the tweets per population hotspots.

Interestingly, neither the Mission Bay nor the Haight-Ashbury neighborhoods are the focal for the coldspots represented in Figure 10 – the tweets showing fear per total tweets in the corpus. The Tenderloin neighborhood shows the most coverage of the 99% confidence

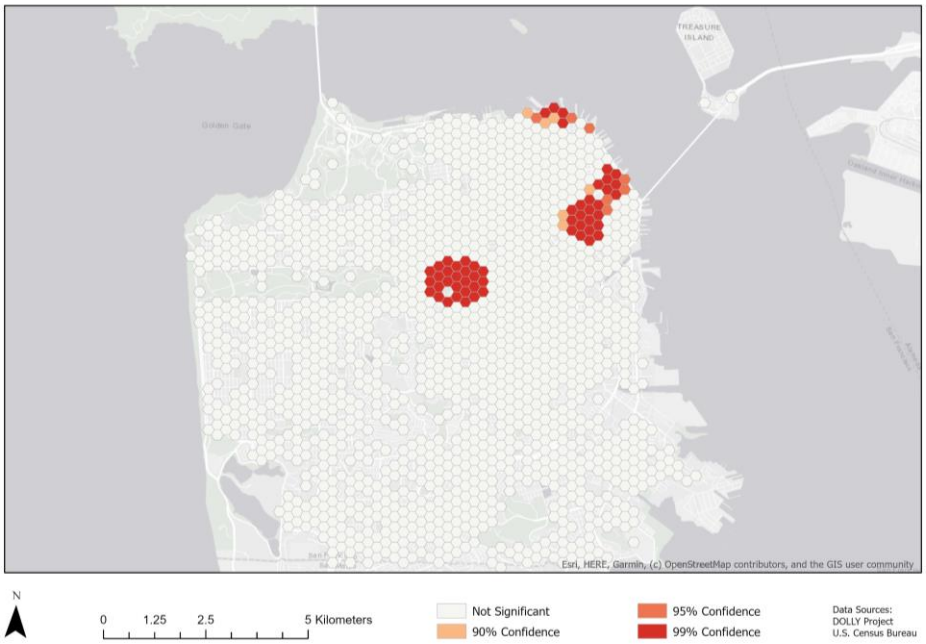

Fig. 7. All Tweets Per Population Hotspots Map – San Francisco areas with different degrees of confidence for hotspots of all tweets in the corpus per population.

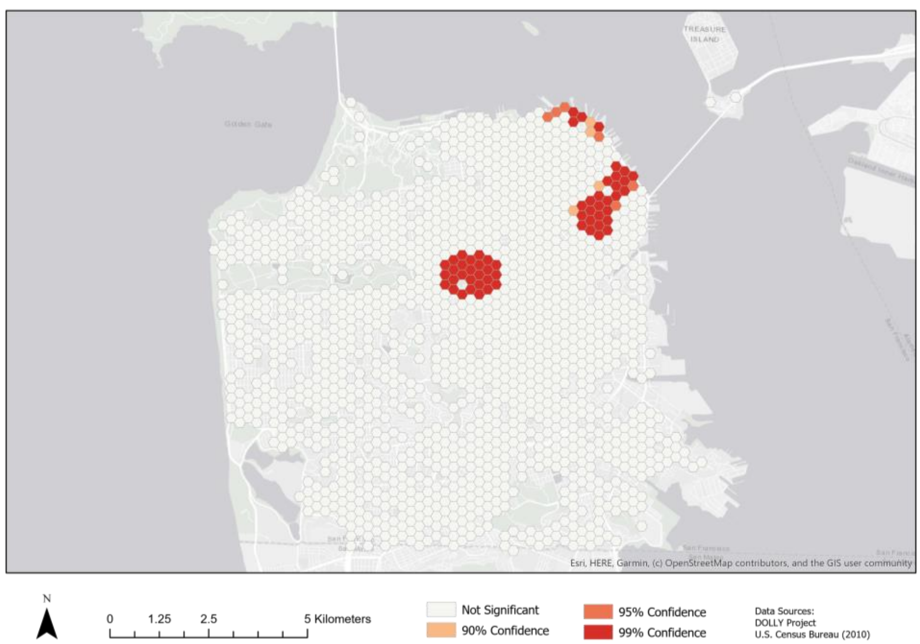

Fig. 8. Tweets Showing Fear Per Population Hotspots Map – San Francisco areas with different degrees of confidence for hotspots of tweets showing fear per population.

coldspots. There is one 90% confidence hotspot in the Oceanview/Merced/Ingleside neighborhood, but the focus for the coldspots is the northeastern areas of the city.

Overall, there does not appear to be any significant relationship between all the tweets in the corpus and tweets showing fear in a temporal sense. Additionally, both datasets do not show any significant relationship in space throughout the city. The coldspots of tweets and tweets showing fear in Figure 10 reinforces the idea that there is no real discerning relationship between the two.

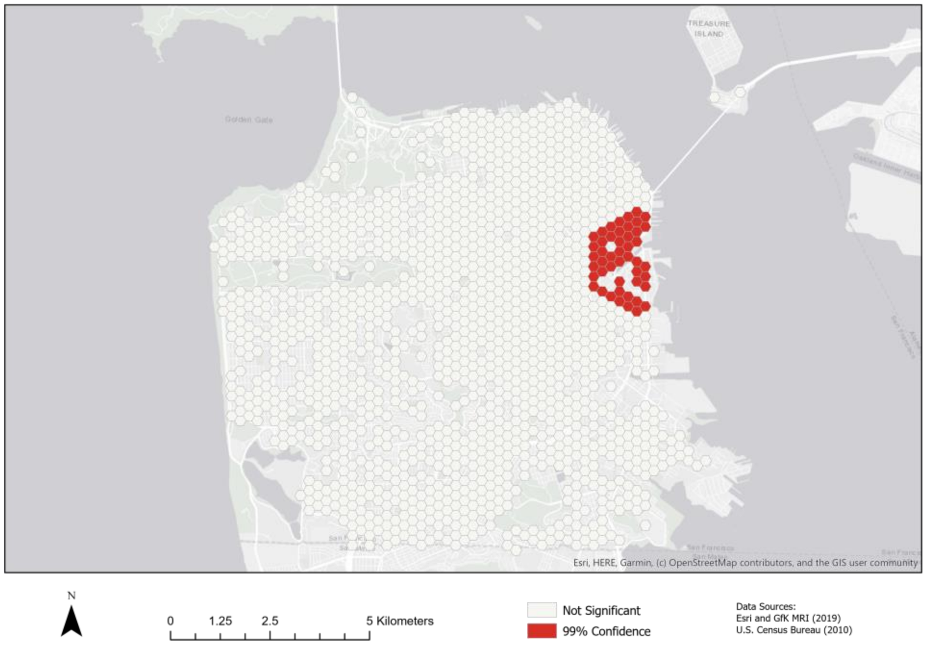

Fig. 9. Smartphones Per Population Hotspots Map – San Francisco areas with different degrees of confidence for hotspots of smartphones per population.

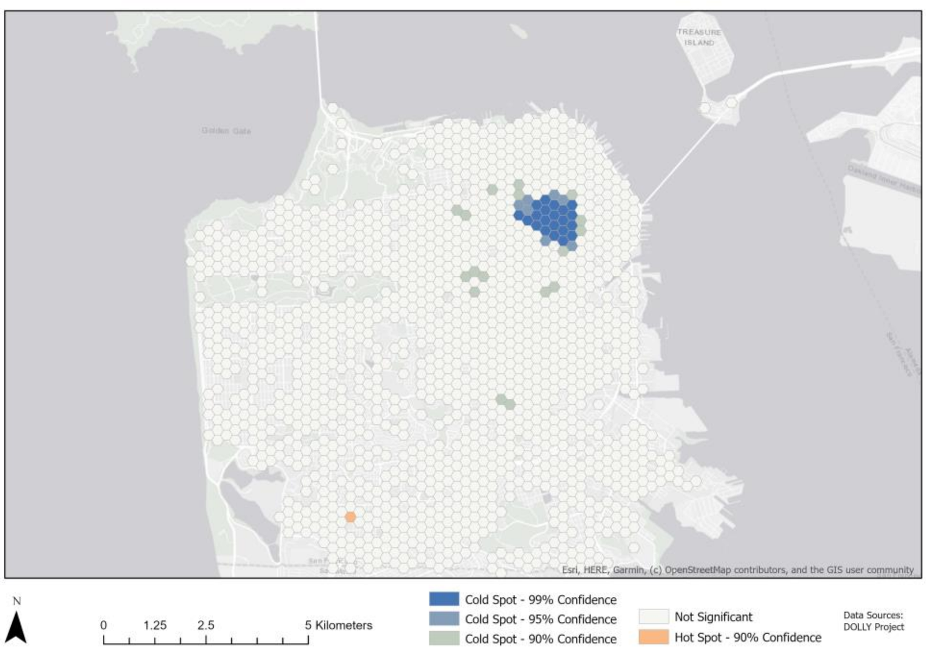

Fig. 10. Tweets Showing Fear Per Total Tweets Hotspots and Coldspots Map – San Francisco areas with different degrees of confidence for hotspots and coldspots of tweets showing fear per total tweets in the corpus.

B. Crime

To understand the broad view of crime in San Francisco and how it has progressed, a look at crime over the course of the previous decade is shown for all crime and Part I crime (see Figure 11). Both record their minimum at the very start of the dataset in January 2010 (12,101 for all crime and 7,151 for Part I crime) and both show a maximum in January 2018 (28,885 for all crime and 21,555 for Part I crime). While there are differences between the peaks and troughs, the general shape between all crime and Part I crime totals remain consistent. Within the full years constrained by the Twitter corpus (2013 through 2017), the year with the lowest totals for all crime and Part I crime is 2013 and the most is 2017. In 2013 the total was 88,685 while in 2017 the total was 98,433, which is an 11% increase. For reference, San Francisco's total population grew from 805,235 in April 2010 to 881,549 in July 2019, a 9.5% increase (United States Census Bureau, n.d.). There is no significant difference in the patterns seen between all crimes and Part I crimes over the time scale displayed in Figure 11.

Boxplots with means for all crime and Part I crime per day of the week are displayed in see Figure 12. The day with the least recorded crime is Sunday for all crime (12,099 counts) and Monday for Part I crime (8,453 counts). Friday records the most crime for both (13,698 for all crime and 8,823 for Part I crimes). No significant difference is seen in the patterns between all crimes and Part I crimes per day of the week.

Continuing down the temporal vein of crime, both all crime and Part I crime show similar trends over the course of a given day (see Figure 13). Both show an absolute minimum at hour 5 (6,249 for all crime and 4,308 for Part I crime), and both show an absolute maximum at hour 18 (39,680 counts for all crime and 30,235 counts for Part I crime). Overall, recorded crime incidents decrease from hour 0 to hour 5 and then increase again until hour 18 at which time there is another decrease until hour 23. The patterns between all crimes and Part I crimes were not significantly different in Figure 13.

The five most common crime incidents recorded in the past decade are shown in Figure 14. In the past decade, the top five categorical crime incidents in San Francisco are larceny,

Fig. 11. Crime in the Past Decade – The total reported, geolocated, and timestamped crimes in San Francisco from the previous decade. Each year marker falls on January 1st. There were only 106 reported crimes from November 1st, 2019 to January 1st, 2020, which was low compared to the prior two-month time step's 14,723 reported crimes, hence the steep drop-off. The data may have been moved to a new system or repository for the next decade.

vandalism, battery, burglary, and auto theft, in that order. Larceny is by far the most frequent crimes committed in San Francisco in the past decade with 387,830 incidents out of the 906,703 crime incidents recorded in the data set that was retrieved from the DataSF repository. Larceny comprises 42.8% of all crimes recorded in the past decade and outnumbers the next highest, vandalism, by nearly a quarter of a million counts; the other four categories swapped rank at certain intervals throughout the decade. Larceny was at a minimum in March 2010 when it comprised 30.1% of crimes committed and was at a maximum in September 2019, constituting 52.2%.

Fig. 12. Crime Per Day of the Week – The means for the reported, geolocated, and timestamped crimes in San Francisco for each day of the week correlating with the tweets' date range.

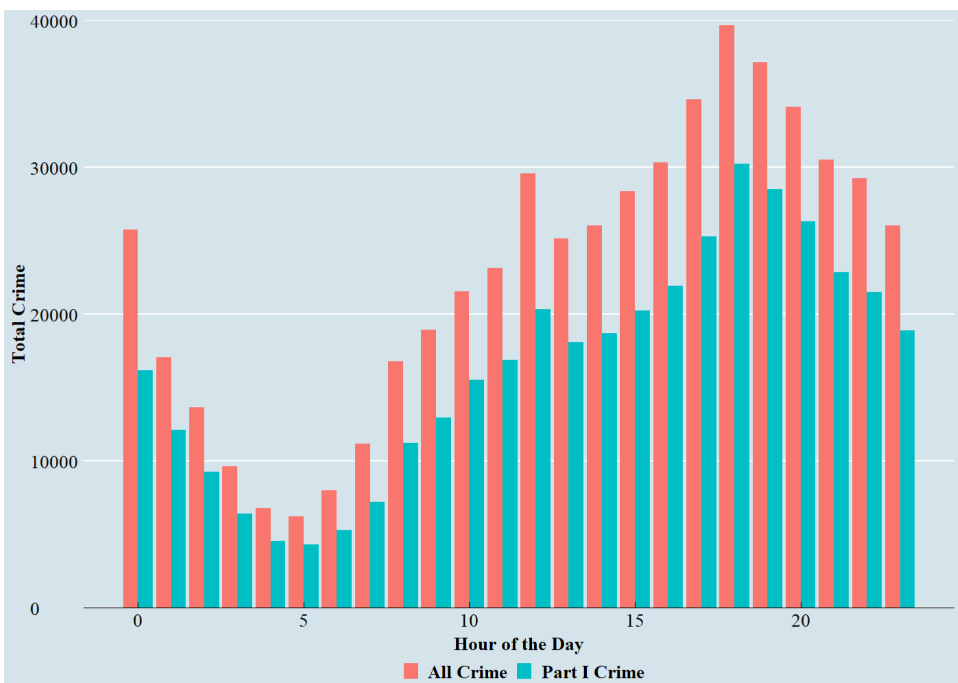

Fig. 13. Crime Per Hour of the Day – The total reported, geolocated, and timestamped crimes in San Francisco for each hour of the day correlating with the tweets' date range.

A view of the ten most frequent crimes and their totals per year is shown in Figure 15. Larceny remains the most frequent crime, and assault, generally is the lowest. Vandalism consistently remains in the top three. For Part I crime, larceny and aggravated assault remain in the top two categories of crime from 2012 through 2018, with homicide the least recorded. Note that aggravated assault is equivalent to aggravated assault *and* battery.

Point density mapping is useful to show the areas in the city that are centers for recorded crime incidents. The point densities for all crime and for Part I crime from 2003 to 2020 show the same general pattern with the majority of the crime incidents concentrated in the northeastern areas, specifically the Tenderloin neighborhood (see Figure 16). There are several other spots of concentrated crime such the Bayview Hunters Point, Mission, and Marina neighborhoods. The major areas for high concentrations of point densities are situated near the intersection of the Tenderloin, South of Market, Financial District/South Beach neighborhoods boundaries.

As previously mentioned, all time step intervals span one month; the last time step interval concludes at the end of March 2018 to align with the date range of the Twitter data.

The hotspots and coldspots for all crime from 2003 to 2020 (not visually depicted in this research) show a dichotomy of two major regions, which split the city from northwest to southeast. The northeastern areas were mostly hotspots while the southwestern areas were mostly coldspots. Interestingly, the area of the city in which there is a high point density of crime in and around the Tenderloin neighborhood shows a center of diminishing hotspots surrounded by intensifying and persistent hotspots. Diminishing hotspots have been statistically significant hotspots for 90% of the time step intervals, which includes the final time step, but overall have statistically significant decreases in intensity of clustering in each time step. The Tenderloin neighborhood shows a relatively high density of crime based on

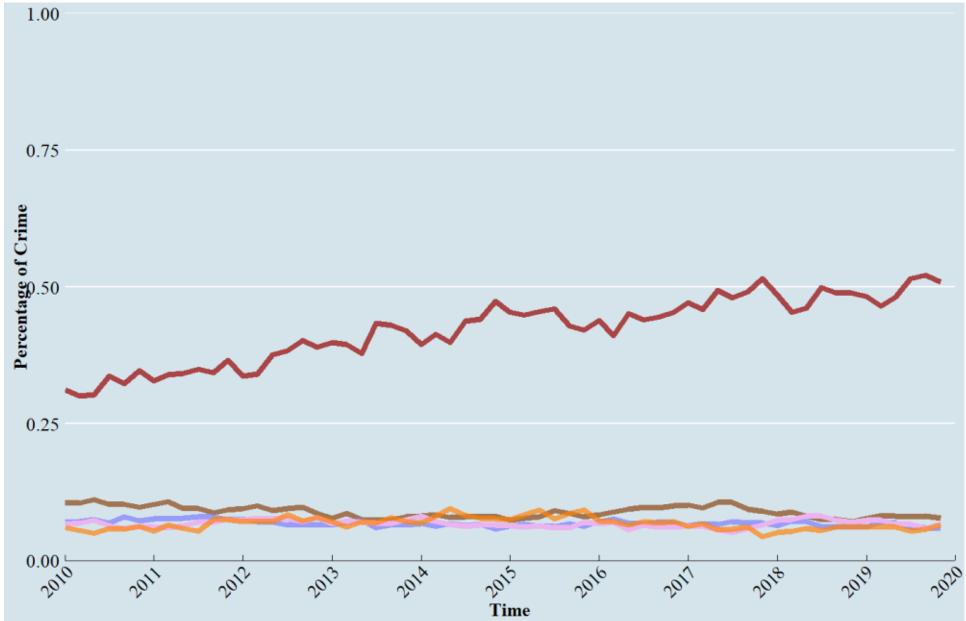

Fig. 14. Top Five Crimes in the Past Decade – The percentage of total reported, geolocated, and timestamped top five crimes in San Francisco from the previous decade. Each year marker falls on January 1st.

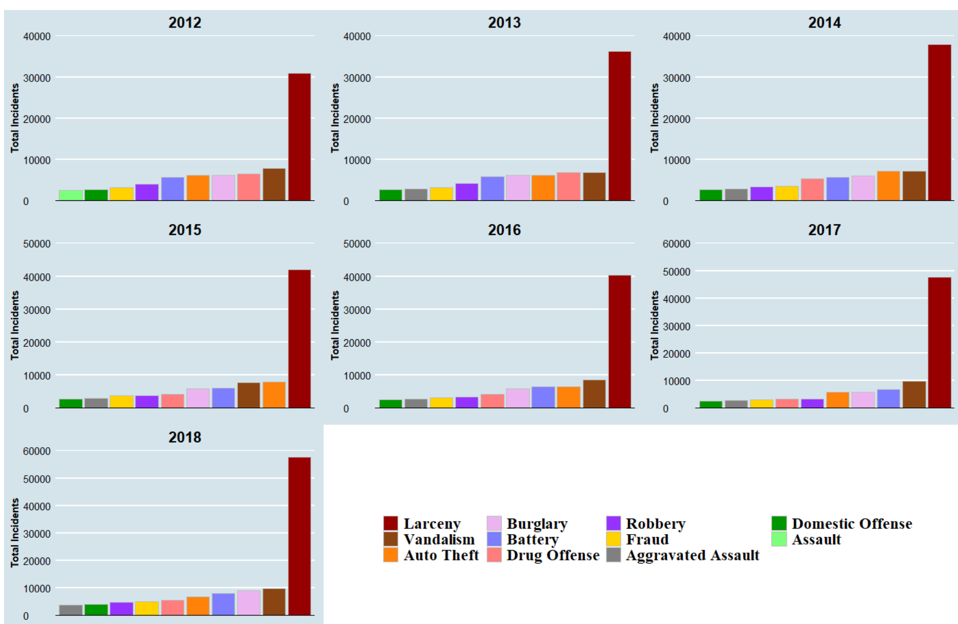

Fig. 15. Top Ten Crimes Per Year – The total reported, geolocated, and timestamped top ten crimes in San Francisco from the 2012 through 2018.

previously reviewed figures, but the diminishing hotspots indicate there was a decreasing trend in the reported crimes progressing from the beginning of the time range corresponding with the Twitter data to the end. However, the areas contiguous with the diminishing hotspots in the Tenderloin neighborhood show intensifying hotspots that have been statistically significant for 90% of the time step intervals and were characterized by a statistically significant increase in intensity of clustering over the entire time range. Also, surrounding the intensifying hotspots in the Tenderloin neighborhood, persistent hotspots have been statistically significant hotspots for 90% of the time step intervals with no trend in increased or decreased intensity of clustering over time. Still in the northeastern areas, the sections of the city closest to the San Francisco-Oakland Bay Bridge have sporadic and oscillating hotspots. Oscillating hotspots, much like sporadic hotspots, have less than 90% of the time step intervals characterized by statistically significant hotspots, however, the final time step interval was a statistically significant hotspot; oscillating hotspots have a history as a statistically significant coldspot in a prior time step. Southwestern areas of the city, the other component of the aforementioned dichotomy, show mostly sporadic coldspots with some intensifying coldspots. The coldspots' categories descriptions are the reverse of the descriptions for hotspots. The Sunset/Parkside neighborhood was mostly characterized by intensifying coldspots while the southern border of the city was mostly sporadic coldspots.

In slight contrast, the hotspots and coldspots for Part I crime from 2003 to 2020 (not depicted in this research) appear more generalized than seen for all crime. Again, the time step interval was one month, and the time range is the full range available for the dataset. It represents a narrower range of hotspot/coldspot categories. The Tenderloin neighborhood in which there are high point densities of crime shows exclusively intensifying hotspots with sporadic hotspots near the San Francisco-Oakland Bay Bridge. The southwestern areas of

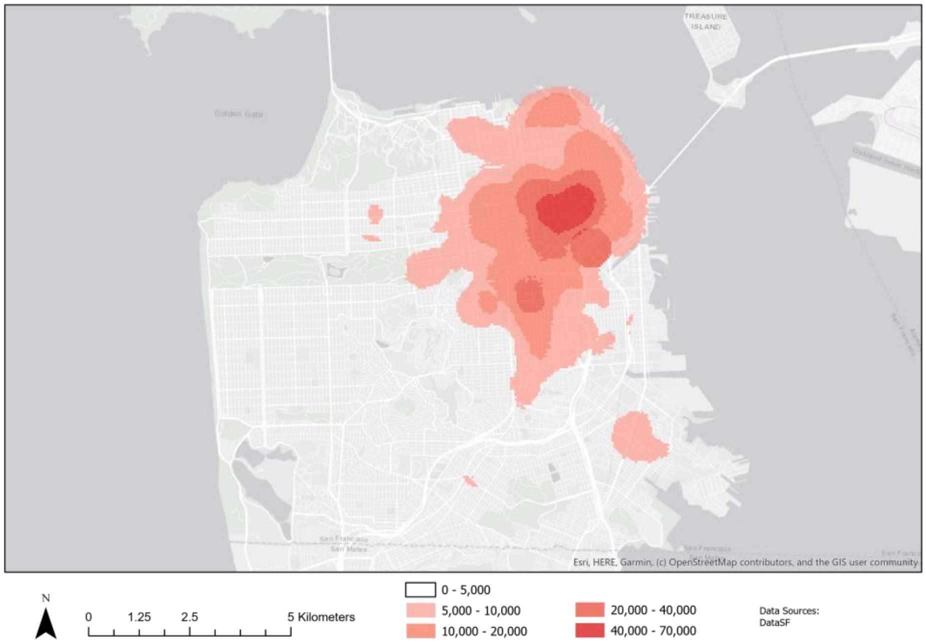

Fig. 16. Crimes from 2012 through 2018 Point Density Map – San Francisco areas with low to relatively high crime incidents corresponding with the tweets' corpus time range. Note the artificial circles due to the radii distance interpolation method.

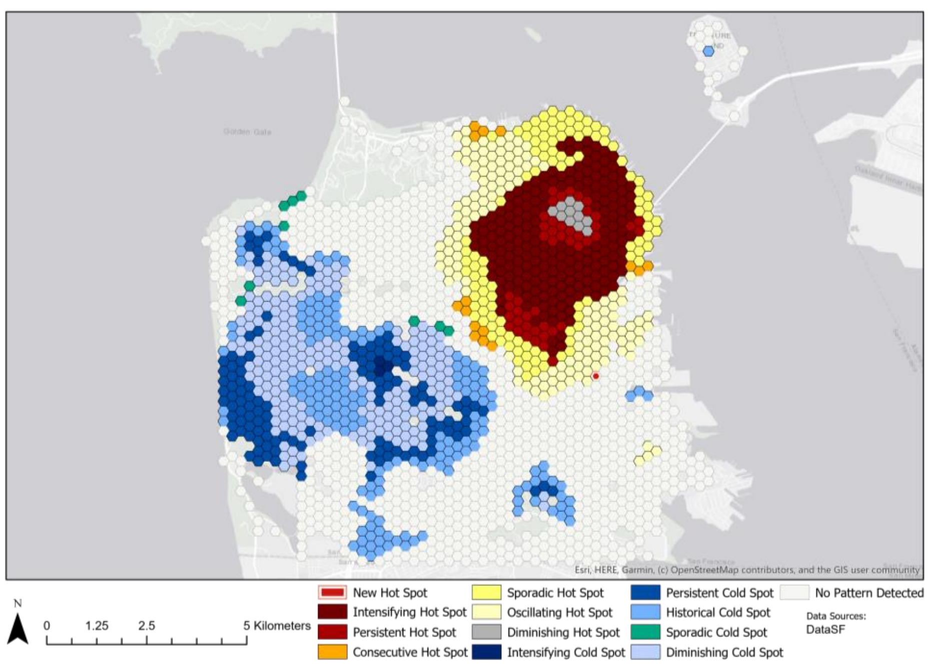

Fig. 17. Crimes from 2012 through 2018 Hotspots and Coldspots Map – San Francisco areas with varying categories of hot/coldspots for crime corresponding with the tweets' corpus time range.

the city show similar patterns to all crime except there were more areas of no detectable hotspot/coldspot patterns.

All crime from 2012 through 2018 show a more detailed view of the emergence of hotspots and coldspots in San Francisco during the years corresponding with the Twitter dataset (see Figure 17). The diminishing hotspots remain in and around the Tenderloin neighborhood, in the northeast with the intensifying and persistent hotspots surrounding it. The historical hotspots near the periphery have been replaced with consecutive hotspots. Historical hotspots are those that have the most recent time step interval characterized as not hot; they do have at least 90% of the time step intervals as statistically significant hotspots. The sporadic hotspots are more prevalent near the San Francisco-Oakland Bay Bridge, and there is a single new hotspot near the intersection of the Potrero Hill, Bayview Hunters Point, Bernal Heights, and Mission neighborhoods. New hotspots are those locations that are statistically significant hotspots for the final time step, but have not been a statistically significant hotspot in any previous time steps. In terms of coldspots, the southern boundary of the city shows a drastic change from sporadic coldspots to areas with no detectable hotspot/coldspot patterns and a subdued presence of intensifying coldspots, which have been replaced with persistent and diminishing coldspots. The hotspots and coldspots for Part I crime from 2012 through 2018 are similar to all crime, except for the replacement of the diminishing hotspots by persistent hotspots.

Instead of simply looking at the hotspots for the crime totals in San Francisco, it is prudent to also view crime against the population as a rate. From 2003 to 2020, all crime per population (see Figure 18) and only the Part I crime per population (see Figure 19) are displayed. The major hotspots with 99% confidence for all crime per population were in the Mission and the West of Twin Peaks neighborhood areas, with small clusters of 90 – 99%

Fig. 18. Crime Per Population Hotspots Map – San Francisco areas with varying confidence of hotspots for crime across the full available crime data set from 2003 to 2020.

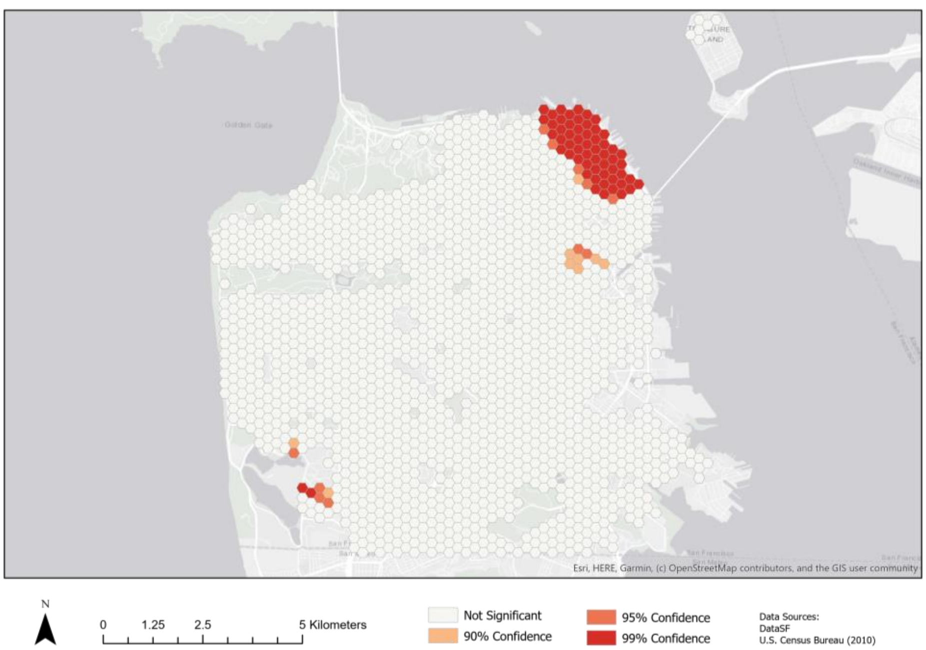

Fig. 19. Part I Crime Per Population Hotspots Map – San Francisco areas with varying confidence of hotspots for Part I crime across the full available crime data set from 2003 to 2020.

confidence hotspots near the North Beach neighborhood area and the adjacent San Francisco-Oakland Bay Bridge. In contrast to Figure 17 for the hotspots and coldspots of crime totals, there is a large cluster of 99% confidence hotspots in the same areas that were categorized as consecutive or diminishing coldspots.

Part I crime per population show a very different depiction from all crime per population. There is a large section in the northeastern areas of the city, spanning the North Beach, Chinatown, and Financial District/South Beach neighborhoods, which contain hotspots with 99% confidence. There is a smaller area of 90% and 95% confidence hotspots in the South of Market neighborhood. Similar to all crime per population, there are clusters of hotspots ranging from 90 – 99% confidence in areas of the southwestern region that were categorized in Figure 15 as sporadic coldspots in the Lakeshore neighborhood.

From 2012 through 2018, all crime per population and Part I crime per population show nearly identical depictions to Part I crime per population from 2003 to 2020 in Figure 19.

Temporally, there is no significant difference between Part I crime and all crime in the city as shown in Figures 11 through 13. Also, as with the Twitter figures, there did not appear to be significant relationships between Part I crime and all crime, spatially, from the point density and emerging hotspots maps shown in Figures 14 through 17. The most significant difference between Part I crime and all crime in the city came when the population was used as a rate in Figures 18 and 19 – minimal overlap between the two.

C. Relationships

The exact spatial and temporal relationships between crime and the tweets showing fear can be visualized and statistically quantified with certain tests. Spatial autocorrelation is one such test. More specifically Global Moran's *I* is a global statistic that can output a single value for an entire dataset comprising crimes and tweets, which can report if variables of interest follow the so-called first law of geography. This first law of geography states that everything is related, however, things that are spatially closer together are more related than things that are further away (Miller, 2004). Explicitly, Global Moran's I compares similarities between objects – hexagons in this case – to their neighbors; afterwards, Global Moran's I averages the similarities and outputs an overall sense of the variables' spatial pattern (Getis & Ord, 1992). To this end, the one-month time step hexagon bins that were created for past emerging hotspot analyses were concatenated to include both datasets for all crime and tweets showing fear. Figure 20 shows the Global Moran's *I* report for the tweets showing fear in San Francisco. The clustering result was the same for the Global Moran's *I*

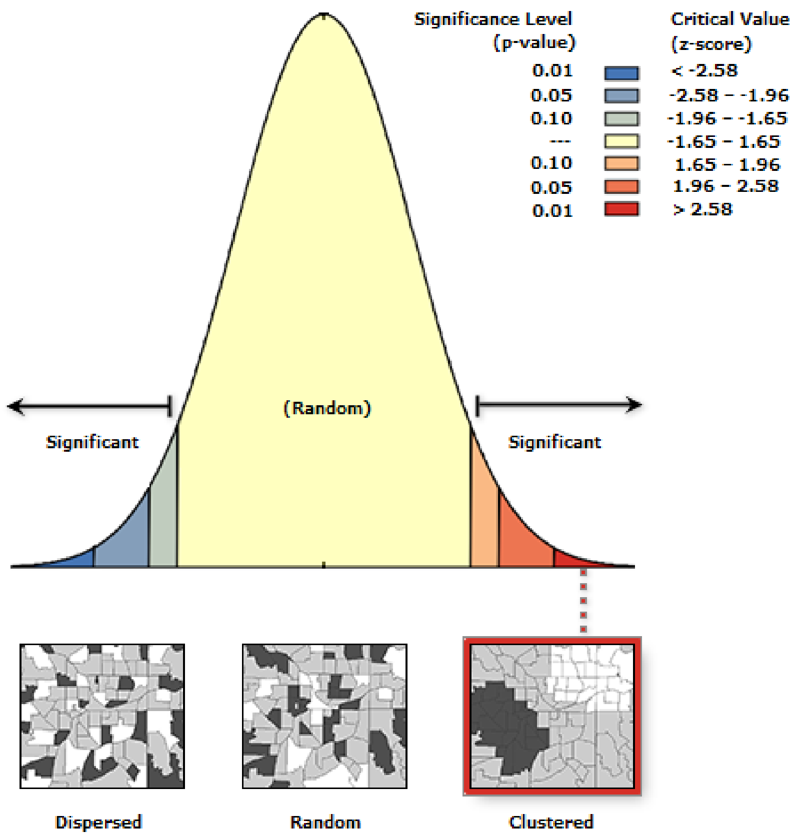

Fig. 20. Tweets' Global Moran's *I* Report – The z-score of 46.53277 provided by the spatial autocorrelation report indicates that there is a less than 1% likelihood that the result of clustered patterning could be random chance (Image source: ArcGIS Pro).

conducted for the crime variable within the same hexagon bins. For both the crime and tweets showing fear variables, the parameters supplied for the spatial autocorrelation were the same: an inverse distance squared conceptualization of spatial relationships, and a Euclidean distance method. The inverse distance squared method was chosen so that only hexagon closest to one another would influence the calculations significantly. The Euclidean distance method was chosen for simplicity. The results for crime indicated a clustered relationship; the results for tweets showing fear indicated a clustered relationship as well.

However, in some cases such as this research, it is more beneficial to understand the relationship between variables. The comparison between two variables is made possible with Local Moran's *I*. The resulting map for Local Moran's I in San Francisco when considering crime and tweets within the same range of time can be seen in Figure 21. The parameters supplied for Local Moran's *I* were the same as the Global Moran's *I*, except an additional parameter was required for the number of permutations conducted on each hexagon bin feature. 499 permutations were chosen to allow for a smaller possible p-value in the calculations. The categorical characterizations for the hexagonal bins were "high-high cluster", "high-low outlier", "low-high outlier", "low-low cluster", and "not significant"; respectively, the categories signify a positive linear correlation between crime and tweets showing fear, a high amount of crime where there is a low amount of tweets showing fear, a low amount of crime where there is a high amount of tweets showing fear, and a negative linear correlation. Based on the map depicted in Figure 21, the conclusion is that the northeastern areas were characterized as having a positive linear relationship between crime and tweets showing fear. This conclusion means that as recorded crime increases so do tweets showing fear. In contrast, the southwestern areas of the city showed a negative linear relationship, which means that as recorded crimes increased tweets showing fear decreased.

As shown in the figures depicting the city spatially in the past section, there were high point densities and hotspots of both tweets and crime in the northeastern areas. The figures and their depictions make the conclusion that there is a positive linear relationship between crime and tweets foreseeable.

Another similar tool to test the spatial relationship between crimes and tweets is the "Local Bivariate Relationships" tool in ArcGIS Pro (see Figure 22). The characterizations in Figure 22 have the added capacity of a 95% level of confidence for the spatial correlations depicted. The central and northeastern areas of the city were regions with 95% confidence of a positive correlation and most of the rest of the map was an "undefined complex", which means that the variables are significantly related, but the relationship cannot be categorized. Lastly, the northeastern and southwestern areas of the city were categorized by a "concave" relationship, which means that on an "x,y" plot the dependent variable (tweets showing fear) progressed as a concave curve as the explanatory variable (crime) increased. For the bivariate relationship depicted in Figure 22, the mean p-value between the two variables was 0.0419.

The separate view of crimes and of tweets did not illuminate any intrinsic characteristics within the datasets, however, when crimes and tweets are compared to each other, spatial relationships arise. The northeastern areas of the city were again represented by the 'high-high' cluster shown around the Tenderloin neighborhood, but the areas of positive linearity between the two datasets shifted towards the north, west, and south. In this case, central areas of the city were shown to be positive linear relationships between crime and tweets, which were not significantly represented beforehand.

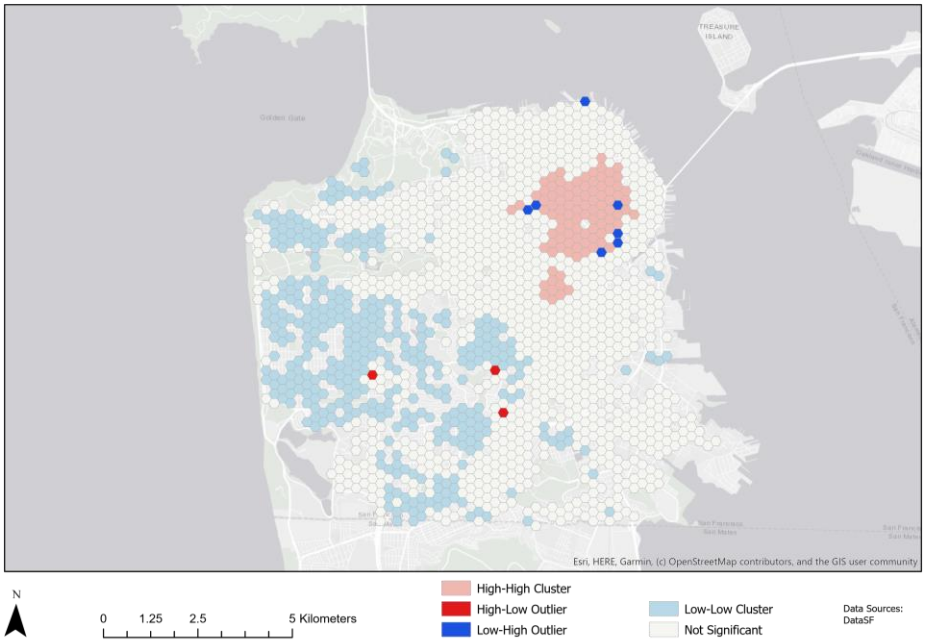

Fig. 21. Crimes and Tweets' Local Moran's *I* Map – San Francisco areas with varying correlation clusters.

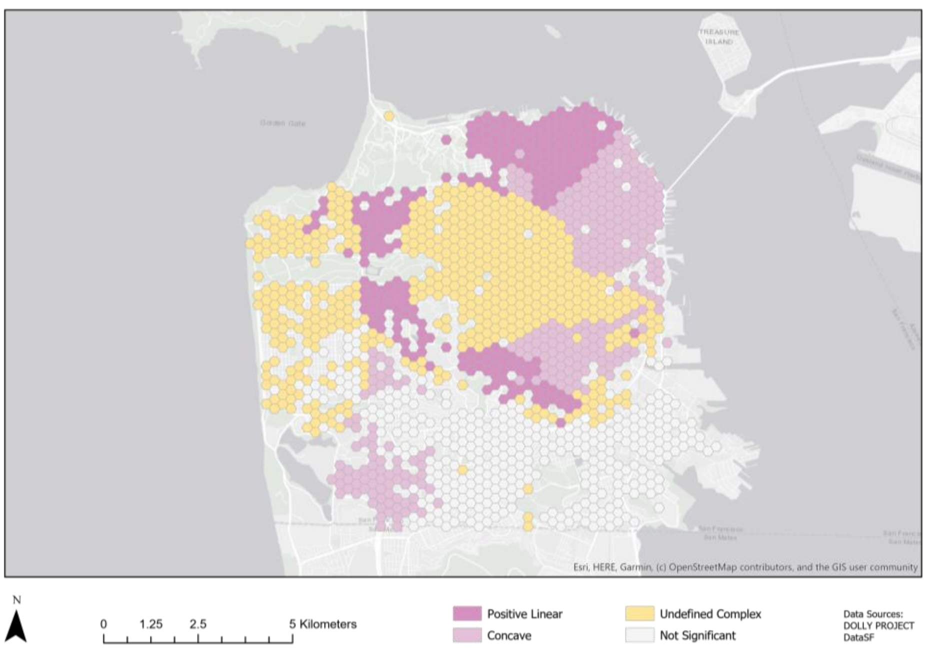

Fig. 22. Crimes and Tweets' Local Bivariate Relationships Map – San Francisco areas with varying bivariate relationships.

VI. Discussion

A. Limitations and Precautions

There are several limitations when conducting a study related to sentiment analysis and crime incidents. For example, 30% of tweets posted are retweets and over 20% of tweets are generated by robots or media tools (Tsou et al., 2013). However, in the current research precautionary measures were taken to correct for this; maintaining retweets or tweets generated by robots within the corpus of this research would have resulted in inaccurate spatial and temporal representations and therefore would have negatively affected the plots and visualizations of the map products.

There is also the issue of demographic representation on Twitter as the users tend to be on average younger and white, which is an issue for measuring population sentiment as it cannot be substituted accurately for a real poll (Diakopoulos & Shamma, 2010).

In terms of the keywords used to collect the tweets, it might have been more prudent to use newspaper sources from San Francisco instead of sources from Los Angeles. However, the degree of error that might have accumulated is uncertain.

In this study, the analytical method used for sentiment analysis was another potential source of error. Previous research has shown that supervised methods such as machine learning are more accurate than unsupervised methods (Bermingham and Smeaton, 2010). However, supervised methods such as machine learning are cost/time intensive.

For the methods used on the tweets, another source of error might have been the random tie-breaker function. The output of tweets categorized with the emotion of fear is slightly different with each iteration as the method of tie breaking is in fact random. This function is statistically sound to avoid bias, but it is worth noting that there are likely small differences in the corpus of tweets used in the statistical analysis and maps. However, each map represented in this study maintains the same corpus of tweets based on the same iteration of tie breaking.

It is worth noting that the following crime analysis should not be considered definitive for all crime incidents that occurred in San Francisco during the time interval characterized by the data. First, the crime incident data is only for those crime incidents that are ascribed a geolocation at the time of the police officer's report. This means that crimes that are not geolocated are not captured in the following analysis. Second, it is unreasonable to assume that every crime incident that occurred in San Francisco would have been reported or cataloged by law enforcement. Essentially, this research should be considered effective for the general relationship between tweets showing fear and crime.

B. Overarching Findings

All crime and Part I crime have absolute maxima situated close to January 2018. For the overall tweets and the tweets showing fear, there are peaks in 2016 following a period of oscillation. However, the absolute maxima for tweets falls roughly 18 months prior to the maxima observed in crimes. The totals per day of the week for crime and for tweets show the same absolute maximum on Friday. The absolute minima and maxima are similar with Sunday or Monday the lowest and Friday the highest in terms of average counts. Tweets showing fear compared to crime presents the interpretation that in San Francisco, Fridays at hour 18 (6 p.m.) correspond to the most documented crime and fear according to Twitter. Temporally, these are similarities that were made apparent through this research.

Point density maps for crime and for tweets expose similar spatial patterns with the densest areas in the northeastern regions of the city. Hotspots maps for crime and for tweets do not show similar results. The hotspots for crime focus around the northeastern areas of the city (Tenderloin neighborhood) while the hotspots for the tweets focus around the central areas of the city (Haight-Ashbury neighborhood). The coldspots for crime are focused around the southwestern areas of the city while there are no apparent coldspots for the tweets.

Comparison between the hotspots maps for the rates of crime and for the rates of tweets reveal discrepancies based upon the time range. Tweets showing fear per population show a more western offset from the crime hotspots per population, particularly for the cluster of 99% confidence hotspots in the Haight-Ashbury neighborhood, which is situated more easterly, in and around the Mission neighborhood (crime from 2003 to 2020). However, the periphery of hotspots near the San Francisco-Oakland Bay Bridge remain similar in both map views of tweets showing fear and of crime. The cluster of hotspots near the eastern bridge remain similar for tweets showing fear per population and crime. In space, there is support for a relationship between crime and fear as garnered from tweets based on the 'high-high' cluster in the northeastern areas and the positive linear relationship in the areas that overlap. The 'high-high' cluster denoting correlation between tweets showing fear and crime resonant with the picture afforded in the separate point density maps made for both the tweets and for crime – the clusters are focused in the same northeastern regions of the city. The areas of the city that show positive linear relationships between tweets showing fear and crime display a different depiction. The positive linear relationships are shown to be more northerly than point density maps hinted and even painted central areas of the city that were not previously represented on other maps in this research.

The spatial similarities between tweets showing fear and crime and the overlap, temporally, for the days of the week and hours of the day support the hypothesis that the two are related. Using large time scales such as the entire previous decade of crime and overlapping that with the tweets is not a good way to delve into the degree to which the two are related. Instead, apparent similarities arose at the weekday and hourly level. Spatially, the same can be said because the north and eastern areas of San Francisco consistently were represented in the maps in this research, however, future research may aim to peer into a larger geographic scale to see if the similarities hold true even down to the street or block level.

Tweets may very well serve as a useful tool in crime studies. This research sought to understand if there were any similarities in time and space between tweets and crime, but it may be prudent for future research to explore this idea more thoroughly by building a model to predict crime based on tweets showing fear. The two datasets aligned more closely at the weekday and hourly scale, which means that there may be some efficacy in using Twitter as a proxy for crime over time. Spatial overlaps and positive linearity in San Francisco was kept to a fairly broad scale, but the similarities found there may warrant future research to study the city in a deeper sense – at the street-level. If the tweets showing fear can serve as a basis for understanding crime, then a predictive model should be the next goal to understand the level of accuracy that can be achieved.

The degree to which tweets can serve as a proxy for crime will be fully understood when predictive models are developed in future research. For now, the current research scratched the surface to understand if predictive models may be a good aim, and based on the relationships in space in the northeastern areas of the city and in time (daily and hourly) it may be fruitful to do so.

VII. Conclusions

This research preprocessed and conducted sentiment analysis on over 215,000 geolocated/timestamped tweets in San Francisco and over 2,500,000 geolocated/timestamped crimes to produce an analysis based on the hypothesis that, in space and time, fear as reflected on Twitter would be highest in areas with the greatest reported crime. Indeed, an interpretation of the results indicated that the northeastern areas of San Francisco were more probable locations for crime, more likely larceny overall or aggravated assault if a Part I crime, on Fridays at 6 p.m., which was reinforced by the tweets showing fear. Unfortunately, if one does become a victim of larceny, there is a little chance of the materials stolen being recovered because it is the least solved crime in San Francisco — just approximately 5% result in arrests compared to homicide at 93% (San Francisco Police Department, 2020). Future research may seek to understand the results from different or more detailed temporal perspectives and test whether the results transcend city location, size, and demographics.

Through preprocessing and sentiment analysis, a synopsis of totals for crime and tweets showing fear within the same time range was produced. All crime and Part I crime were at a minimum in 2013 and a maximum in 2018 while tweets were at a minimum in 2012 and a maximum in 2016. For all crime and Part I crime, April was the month with the lowest incidents, January the highest, yet January was the lowest for tweets, and March the highest. In a given week, Tuesday was the safest from crime and Friday was the most dangerous in terms of crime totals, which was echoed by the tweets showing fear totals. On a given day, crime incidents decreased from midnight to 5 A.M. and reached their height at 6 P.M., which was again echoed by the tweets except for the minimum, which was 10 A.M.

In this study, a synopsis of spatial patterns for crime and tweets showing fear within the same time range can be reviewed. The point densities for all crime and Part I crime were concentrated in the northeastern areas of the city, greatest in the Tenderloin neighborhood; the tweets were focused on the Haight-Ashbury neighborhood, near the center of the city, with significant hotspot areas in and around the northeastern areas such as the Tenderloin, South of Market, and Financial District/South Beach neighborhoods. Hotspots for all crime and Part I crime were focused in and around the Tenderloin neighborhood in the northeast, while hotspots for tweets were situated in and around the Haight-Ashbury neighborhood. However, there was overlap of the consecutive and sporadic hotspots for both crime and tweets. All crime and Part I crime per population showed the majority of the 99% confidence hotspots in the far northeastern areas of the city, spanning the North Beach, Chinatown, and Financial District/South Beach neighborhoods, and the tweets depicted most of their 99% hotspots in and around the Haight-Ashbury neighborhood as well as the North Beach, Financial District/South Beach, South of Market neighborhoods. The per population hotspots for crime and tweets showed overlaps of 99% confidence in the North Beach and Financial District/South Beach neighborhoods. Later hotspot patterns for crime and tweets showing fear, past the 2018 cutoff, may be similar to those previously displayed because hotspots are intrinsic representations of crime in the future, as they are based on crime in the past (Chainey, Tompson, & Uhlig, 2008).

This research could be beneficial to law enforcement for efficient allocation of resources and manpower in order to curtail crime in the areas that most affect residents' perception of safety. Simple measures could also be taken to ease some fears. More prevalent street lighting can go a long way to counteract the fear of crime (Hale, 1996). However, also according to Hale (1996), there is "no single counter to the fear of crime for every area," which means that countermeasures would have to be robust and versatile for a wider range of environments. Decay and disorder in a living environment can be helped and alleviated by denizens with a proactive rather than reactive willingness to regain control of their sense of safety (Hale, 1996).

References

Aslam, S. (2020, February 10). Twitter by the Numbers (2020): Stats, Demographics & Fun Facts. *Omnicore Agency*. Retrieved May 30, 2020, from https://www.omnicoreagency.com/twitter-statistics/

Asur, S., & Huberman, B. A. (2010). Predicting the Future with Social Media. *2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, 492–499. https://doi.org/10.1109/WI-IAT.2010.63

Birch, C. P. D., Oom, S. P., & Beecham, J. A. (2007). Rectangular and hexagonal grids used for observation, experiment and simulation in ecology. *Ecological Modelling*, *206*(3), 347–359. https://doi.org/10.1016/j.ecolmodel.2007.03.041

Bollen, J., Mao, H., & Zeng, X.-J. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, *2*(1), 1–8. https://doi.org/10.1016/j.jocs.2010.12.007

Box, S., Hale, C., & Andrews, G. (1988). Explaining Fear of Crime. *The British Journal of Criminology*, *28*(3), 340–356. JSTOR.

Chainey, S., Tompson, L., & Uhlig, S. (2008). The Utility of Hotspot Mapping for Predicting Spatial Patterns of Crime. *Security Journal*, *21*(1–2), 4–28. https://doi.org/10.1057/palgrave.sj.8350066

Chaussé, P. (2010). Computing Generalized Method of Moments and Generalized Empirical Likelihood with *R*. *Journal of Statistical Software*, *34*(11). https://doi.org/10.18637/jss.v034.i11

Clearance Rates Dashboard. (2020, April 20). San Francisco Police Department. Retrieved June 2, 2020, from http://www.sanfranciscopolice.org/stay-safe/crime-data/clearance-rates-dashboard

Clement, J. (2019, August 9). *Twitter usage reach in the United States 2019, by age group*. Statista. Retrieved May 29, 2020, from https://www.statista.com/statistics/265647/share-of-us-internet-users-who-use-twitter-by-age-group/

Clement, J. (2020, February 20). *Twitter - Statistics & Facts*. Statista. Retrieved May 2, 2020, from https://www.statista.com/topics/737/twitter/

Diakopoulos, N. A., & Shamma, D. A. (2010). Characterizing debate performance via aggregated twitter sentiment. *Proceedings of the 28th International Conference on Human Factors in Computing Systems - CHI '10*, 1195. https://doi.org/10.1145/1753326.1753504

Ferraro, K. F., & Grange, R. L. (1987). The Measurement of Fear of Crime. *Sociological Inquiry*, *57*(1), 70–97. https://doi.org/10.1111/j.1475-682X.1987.tb01181.x

Getis, A., & Ord, J. K. (1992). The Analysis of Spatial Association by Use of Distance Statistics. *Geographical Analysis*, *24*(3), 189–206. https://doi.org/10.1111/j.1538-4632.1992.tb00261.x

Gonçalves, P., Araújo, M., Benevenuto, F., & Cha, M. (2013). Comparing and combining sentiment analysis methods. *Proceedings of the First ACM Conference on Online Social Networks*, 27–38. https://doi.org/10.1145/2512938.2512951

Hale, C. (1996). Fear of Crime: A Review of the Literature. *International Review of Victimology*, *4*(2), 79–150. https://doi.org/10.1177/026975809600400201

Hatzivassiloglou, V., & McKeown, K. R. (1997). Predicting the semantic orientation of adjectives. *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics and Eighth Conference of the European Chapter of the Association for Computational Linguistics*, 174–181. https://doi.org/10.3115/976909.979640

Hatzivassiloglou, V., & Wiebe, J. M. (2000). Effects of adjective orientation and gradability on sentence subjectivity. *Proceedings of the 18th Conference on Computational Linguistics*, *1*, 299–305. https://doi.org/10.3115/990820.990864

Intravia, J., Wolff, K. T., Paez, R., & Gibbs, B. R. (2017). Investigating the relationship between social media consumption and fear of crime: A partial analysis of mostly young adults. *Computers in Human Behavior*, *77*, 158–168. https://doi.org/10.1016/j.chb.2017.08.047

Liu, B. (2012). Sentiment Analysis and Opinion Mining. *Synthesis Lectures on Human Language Technologies*, *5*(1), 1–167. https://doi.org/10.2200/S00416ED1V01Y201204HLT016)

Miller, H. J. (2004). Tobler's First Law and Spatial Analysis. *Annals of the Association of American Geographers*, *94*(2), 284–289. JSTOR.

Moon, A. (2017, September 8). *Two-thirds of American adults get news from social media: Survey*. Reuters. Retrieved May 30, 2020, from https://www.reuters.com/article/us-usa-internet-socialmedia-idUSKCN1BJ2A8

Pak, A., & Paroubek, P. (2016). Twitter as a Corpus for Sentiment Analysis and Opinion Mining. *IJARCCE*, *5*(12), 320–322. https://doi.org/10.17148/IJARCCE.2016.51274

Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval.* Vol. 2, No 1-2.

Pobuda, M. (n.d.) *Analyze Crime Using Statistics and the R-ArcGIS Bridge | Learn ArcGIS*. Esri. Retrieved June 2, 2020, from https://learn.arcgis.com/en/projects/analyze-crime-using-statistics-and-the-r-arcgis-bridge/

Smith, K. (2020, January 2). *60 Incredible and Interesting Twitter Stats and Statistics*. Brandwatch. Retrieved May 2, 2020, from https://www.brandwatch.com/blog/twitter-stats-and-statistics/

Social Media Fact Sheet. (2019, June 12). Pew Research Center: Internet, Science & Tech. Retrieved August 23, 2020, from https://www.pewresearch.org/internet/fact-sheet/social-media/

Tasci, S., & Gungor, T. (2009). LDA-based keyword selection in text categorization. *2009 24th International Symposium on Computer and Information Sciences*, 230–235. https://doi.org/10.1109/ISCIS.2009.5291818

Top 10 California Daily Newspapers by Circulation. (2015, September 16). Agility PR Solutions. https://www.agilitypr.com/resources/top-media-outlets/top-10-california-daily-newspapers-by-circulation/

Tsou, M.-H., Yang, J.-A., Lusher, D., Han, S., Spitzberg, B., Gawron, J. M., Gupta, D., & An, L. (2013). Mapping social activities and concepts with social media (Twitter) and web search engines (Yahoo and Bing): A case study in 2012 US Presidential Election. *Cartography and Geographic Information Science*, *40*(4), 337–348. https://doi.org/10.1080/15230406.2013.799738

Turney, P. D. (2001). Thumbs up or thumbs down?: Semantic orientation applied to unsupervised classification of reviews. *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics - ACL '02*, 417. https://doi.org/10.3115/1073083.1073153

U.S. Census Bureau QuickFacts: San Francisco County, California. (n.d.). The United States Census Bureau. Retrieved May 22, 2020, from https://www.census.gov/quickfacts/sanfranciscocountycalifornia

Uniform Crime Reporting Handbook. (2004). U.S. Department of Justice, Federal Bureau of Investigation. Retrieved May 30, 2020, from https://ucr.fbi.gov/additional-ucr-publications/ucr_handbook.pdf

Wang, X., Gerber, M., & Brown, D. (2012). Automatic Crime Prediction Using Events Extracted from Twitter. *Social Computing, Behavioral - Cultural Modeling and Prediction: 5th International Conference, SBP 2012, College Park, MD, USA, April 3-5, 2012. Proceedings* (Vol. 7227). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-29047-3

Yates, A., & Goharian, N. (2013). ADRTrace: Detecting Expected and Unexpected Adverse Drug Reactions from User Reviews on Social Media Sites. In *Advances in Information Retrieval* (Vol. 7814, pp. 816–819). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-36973-5_92

Appendices

R coding language:
R Core Team (2020). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/

Source for San Francisco crime datasets, zoning, police stations/districts, and AOV:
Office of the Chief Data Officer | City and County of San Francisco. (n.d.). DataSF. Retrieved June 2, 2020, from https://datasf.org/

Source for San Francisco roads and highways:
TIGER/Line Shapefile, 2017, county, San Francisco County, CA, All Roads County-based Shapefile—Data.gov. (2017). Retrieved June 2, 2020, from https://catalog.data.gov/dataset/tiger-line-shapefile-2017-county-san-francisco-county-ca-all-roads-county-based-shapefile

Source for sentiment code examples in R:
Silge, J. (2015, December 22). *Joy to the World, and also Anticipation, Disgust, Surprise...* Julia Silge. https://juliasilge.com/blog/joy-to-the-world/

Source for Emerging Hotspots code examples in R:
Pobuda, M. (n.d.) *Analyze Crime Using Statistics and the R-ArcGIS Bridge | Learn ArcGIS*. Esri. Retrieved June 2, 2020, from https://learn.arcgis.com/en/projects/analyze-crime-using-statistics-and-the-r-arcgis-bridge/