Tweet Sentiment Analysis with Pronoun Choice Reveals Online Community Dynamics in Response to Crisis Events

4 authors:

Samira Shaikh
University of North Carolina at Charlotte
50 PUBLICATIONS   **198** CITATIONS

SEE PROFILE

Laurie Beth Feldman
University at Albany, The State University of New York
116 PUBLICATIONS   **3,615** CITATIONS

SEE PROFILE

Eliza Barach
University at Albany, The State University of New York
9 PUBLICATIONS   **9** CITATIONS

SEE PROFILE

Yousri Marzouki
Aix-Marseille Université
51 PUBLICATIONS   **219** CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:

Morphological processing View project

Naive Discrimination Learning (NDL) View project

Tweet Sentiment Analysis with Pronoun Choice Reveals Online Community Dynamics in Response to Crisis Events

Samira Shaikh[1], Laurie Beth Feldman[1], Eliza Barach[1],Yousri Marzouki[2]

[1]State University of New York – University at Albany,
1400 Washington Avenue Albany NY 12222
Qatar University and Aix-Marseille University, Marseille, France
{sshaikh@albany.edu, lfeldman@albany.edu, eliza.barach@gmail.com,
yousri.marzouki@gmail.com}

Abstract. We describe the emergence of an online community from naturally occurring social media data. Our method uses patterns of word choice in an online social platform to characterize how a community forms in response to adverse events such as a terrorist attack. Our focus is English Twitter messages after the Charlie Hebdo terrorist attack in Paris in January 2015). We examined the text to find lexical variation associated with measures of valence, arousal and concreteness. We also examine the patterns of language use of the most prolific twitter users (top 2% by number of tweets) and the most frequent tweets in our collection (top 2% by number of retweets). Differences between users and tweets based on frequency are revealing about how lexical variation in tweeting behavior reflects evolution of a community in reaction to crisis events on an international scale.

Keywords: Human Factors; Sentiment Analysis; Social Media; Natural Language; Big Data

1   Introduction

Variety in online communication is now a popular subject of investigation. These new communication formats range from synchronous to asynchronous, from private to public and from spontaneous to triggered by specific events [10]. These communications often have no specific addressee and yet they can be very revealing about the sender and those with whom she/he communicates. Further, because the roles of sender and recipient continually alternate, the distinction between roles typically loses importance. We use patterns of interaction at the level of groups of users rather than individuals to define community and subcommunity and this is our focus of study.

Dubois and Gaffney [7] and Xu et al. [12] showed that opinion leaders and influencers could be quantitatively identified in Twitter-mediated political discussions. Others differentiate between tweeters who are at the epicenter and those who are more peripheral to the focal event and emphasize how the synergy between them scales up the visibility of an event and helps in disseminating information [4]. Giglietto and Lee

[9] report that retweets are particularly characteristic of emergency events when compared to ceremonial events (e.g., royal birth) that elicit more original tweets. Tweet popularity within a set of hashtags provides a means to detect whether tweeters have a bias to expose themselves selectively to those with similar ideological positions or whether the tweeting environment is more consistent with cross-ideological dissemination [3].

Patterns of lexical variation within a community including the use of emoticons intermixed with words within a bilingual work environment have been informative about language background [1]. The choice of personal pronouns in behavioral descriptions collected across diverse corpora has been revealing about group dynamics and psychological states [6]. Sometimes, relative to self-excluding pronouns (HE, SHE, THEY), news reports with self-including pronouns (I, WE) are associated with more positively valenced accounts [8].

Using Twitter data as a corpus to investigate the conversational characteristics of users is a timely scientific endeavor. For example, Burns and Stieglitz [5] highlighted the importance of studying a shared data set that is built upon one hashtag and stated that hashtags can be considered as "shared conversational markers", because they are inclusive without requiring a uniformity of opinion. In addition, they scale up faster than other types of discourse.

The following study aimed at extending the work conducted by Giglietto and Lee [9] about the twitter hashtags as behavioral indicators of online discourse and counter-discourse dynamics when users are facing a crisis event. Giglietto and Lee [9] used network co-occurrence patterns and semantic cluster analyses of content by hashtags to show that tweets associated with the shooting attack by Islamist gunmen at the office of the Charle Hebdo publication served as a genuine vehicle to exhibit various characteristics of the users' online discussions and reactions. Comparative analyses across hashtags #JeSuisCharlie and #JeNeSuisPasCharlie revealed interesting hidden patterns such as activism, grief, resistance, ethnocentrism and islamophobia in the 74,074 posted tweets collected between 7[th] and 11[th] of January 2015.

The current study analyzed a more sizable sample than the one used by Giglietto and Lee [9] and within a more stable time window, covering the period from January 14[th] to February 9[th]. Moreover, we used many diverse hashtags to cover a sample with different emotional reactions. More importantly, instead of tagging the tweets associated with one opposing hashtag, we opted for a more fine-grained analysis of lexical variation, by differentiating among tweets according to pronoun use in natural human discourse. Our goal was to reveal nuanced discrepancies of group identities originating from the target hashtag #JeSuisCharlie. Indeed, the "Je" in the hashtag refers to the first person pronoun in French (I), so tracking the dynamics of the discourse and counter-discourse based on pronouns and the key words with which they occur is particularly illuminating. We have hypothesized that the presence of the "Je" in the key hashtag of the whole event is a trigger and a prime for an identity issue that a close analysis will help us decipher and understand.

2   Data Collection and Method

The data we collected consisted of all the tweets that contained at least one of 16 hashtags pertaining to the Charlie Hebdo attacks collected via Twitter streaming API. The final dataset consisted of 404,918 tweets from about 190,000 unique twitter ids that were shared between January 14[th] and February 9[th] 2015. The unique hashtags across the corpus are listed in Table 1.

Table 1. List of unique hashtags in our tweet collection.

#JeSuisCharlie	#IamCharlie	#MarcheRepublicaine	#JeSuisAhmed
#ParisShooting	#NousSommesCharlie	#LaFranceEstCharlie	#jesuisfranck
#lemondeestcharlie	#ThanksTheWorldFromFrance	#jenesuispascharlie	#parisattacks
#parisattacks	#parisattack	#ParisEstCharlie	#CharlieHebdo

3   Data analysis

To ascertain that lexical variation is associated with pronoun choice, we extracted all tweets with pronouns and classified them as to whether a pronoun was present or absent and then whether the pronoun was in first (I, ME, MINE, etc.), second (YOU, YOUR, etc.) or third (HE, HIM, HIS, SHE, HER, HERS, THEY, etc.) person. For first person pronouns, we separated them into singular and plural so that we could track individual and collective perspectives [8].

In this article, we analyzed three sets of data for valence, arousal and concreteness measures. The first set is all of the tweets containing first-person singular pronouns, the second is the set of all tweets containing first-person plural pronouns and the last is the set of all tweets containing no first- or second- person pronouns. We defined the last set as the baseline set. In future work, we aim to include tweets containing second-person pronouns (YOU, YOUR etc) in our pronoun analyses.

Each of the three sets was pre-processed as follows: We first removed stop-words (non-content words such as THE, A etc.) from the tweets. Next, we converted all words to their lower-case equivalents. In the last step, we calculated the frequencies of all words that co-occur with our target pronoun set (first-person singular, first-person plural and baseline). Thus, we obtained a list of words, ranked by their frequencies that co-occurred with the sets of pronouns we were interested in. We then determined the valence, arousal and concreteness scores of any word that was listed in our corpus by consulting databases collected by researchers [2] [11]. Psycholinguistic measures [2] [11] included valence (the pleasantness of a given word), arousal (intensity of evoked emotion, terrified > grief) and concreteness (the degree to which the concept is perceptible to one of the senses). We calculated weighted means for arousal valence and concreteness for each word using the following formula: multiplying the arousal, valence and concreteness score of a word by the number of times that word appeared in the corpus. When considering the valence, arousal and concreteness

scores, we focused on those words that exhibited the most extreme scores – which are the 25% most negative words and 25% most positive words for valence, and similarly for arousal and concreteness, the top most 25% words from the respective databases were used.

4　Results

We provide descriptive statistics of the data in Section 4.1. In Sections 4.2-4.4, we present the analysis of tweets that contrasts the measures of valence, arousal and concreteness across the three different sets defined by pronoun use (first person singular, first person plural and no first person pronouns: the baseline). In Section 4.5, we contrast the language use patterns for the subcommunity of super-tweeters (top 2%) and the remainder of the community. In Section 4.6, we contrast the language use patterns that characterize the super-tweets (top 2%) vs. the remainder of our tweet collection.

4.1　Distribution of Tweets and Users

Figure 1 shows the distribution of tweets in our data collection. We note that the rate of tweets was higher in the days immediately after the event, than as time passed.
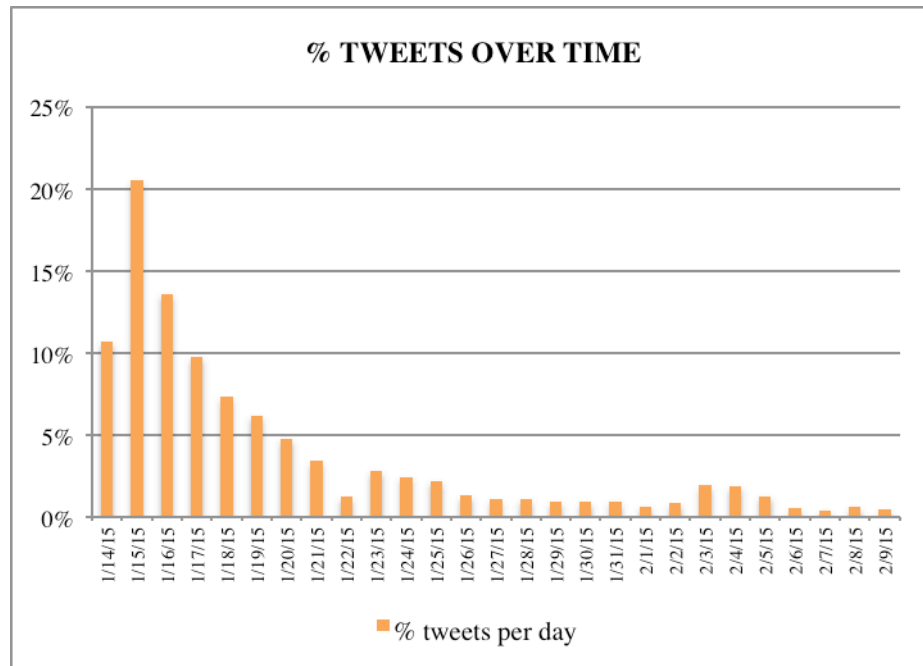

Fig. 1. Histogram of tweets rate over time in our collection

Figure 2 shows the histogram of unique twitter users per day ascertained by the number of unique twitter IDs per day. In Figure 3, we show a word cloud with the most frequently appearing words in our collection.
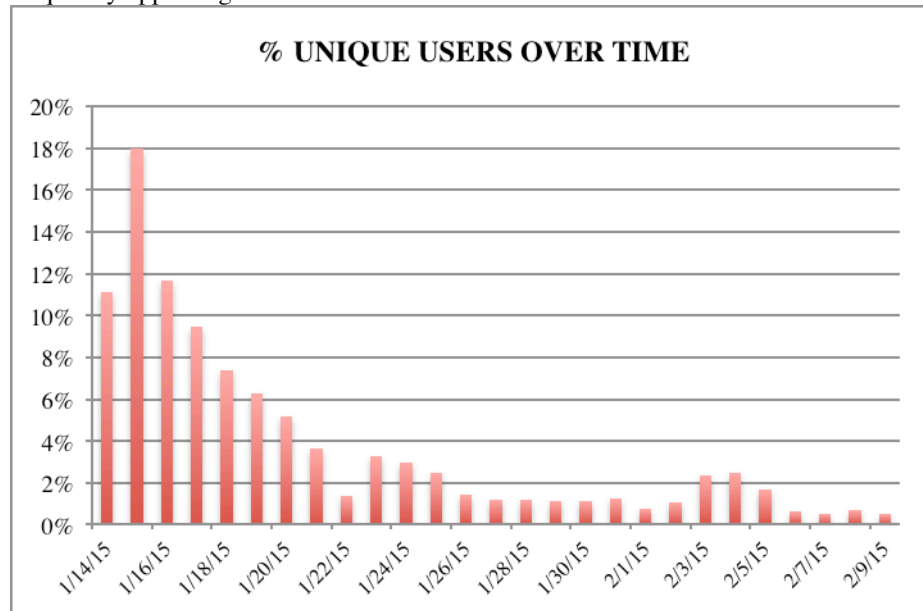

Fig. 2. Histogram of unique twitter users over time

Fig. 3. Word cloud showing the most common words in the collection

4.2 Differences in Valence Scores across Sets of Tweets with Different Pronouns

The analysis across days of the valence of words in tweets containing first-person pronouns indicated that they were more strongly valenced than baseline tweets (without first-person pronouns). The scatter plot and regression lines are shown in Figure 4.

Beta values were -0.0042, 0.0020 and 0.0077 for first-person singular, first-person plural and baseline, respectively. Importantly, slopes are marginally different between first-person singular and first-person plural. Tweets using first-person singular were more negative that tweets using first-person plural and the changes in that difference over time were marginally significant (p = .06537). The presence of strongly valenced language around pronouns suggest greater diffusion of emotion when contrasted with the more impersonal baseline.
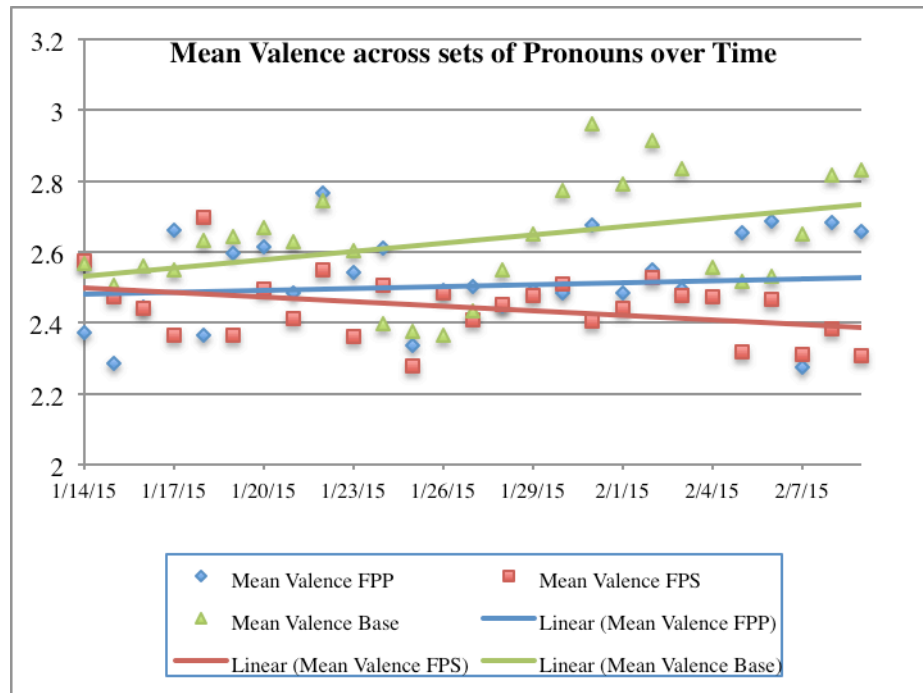

Fig. 4. Scatter diagram with regression lines over mean valence of words that co-occur with first-person singular (FPS) and first-person plural (FFP) pronouns and baseline (no first person pronouns).

4.3 Differences in Arousal Scores across Sets of Tweets with Different Pronouns

The analysis across days of the arousal for words indicated that tweets using first-person singular pronouns tended to have lower arousal than tweets using first-person plural or the baseline tweets. The scatter plot and regression lines are shown in Figure 5. Means were 5.96, 6.20, and 6.15 for first-person singular, first-person plural and baseline respectively. Tweets with first-person singular pronouns failed to show a change in arousal over time. However, tweets using first-person plural pronouns and baseline showed a non-significant decrease in arousal over time (beta = 0.0044, 0.0133 and -0.0245 for first-person singular, first-person plural and baseline respec-

tively; p >.1). This finding is consistent with generally high arousal within the community and minimal change with the passage of time.
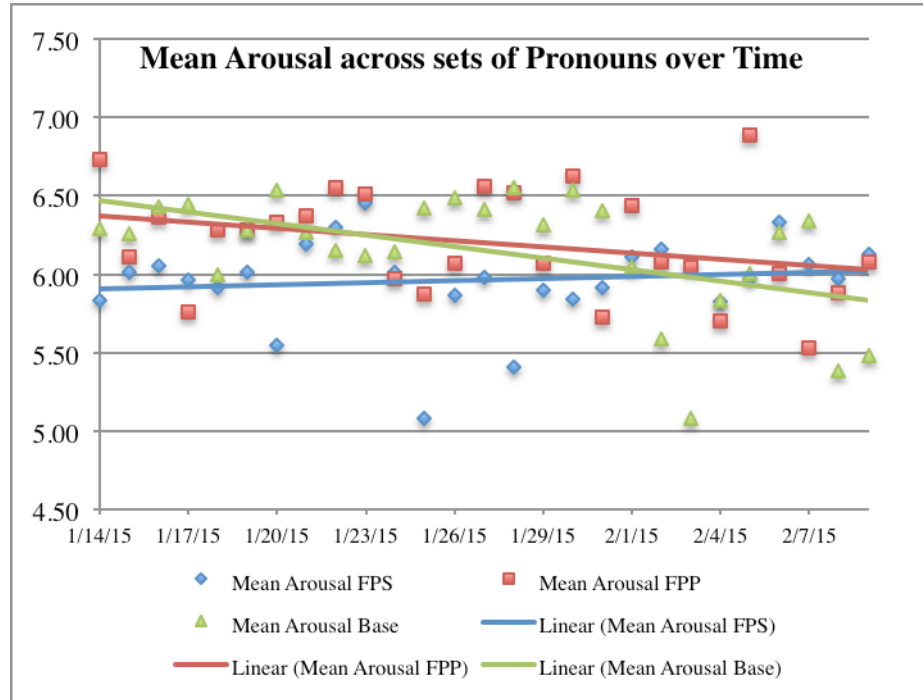

Fig. 5. Scatter diagram with regression lines over mean arousal of words that co-occur with first-person singular (FPS) and first-person plural (FFP) pronouns and baseline (no first person pronouns).

4.4 Differences in Concreteness Scores across Sets of Tweets with Different Pronouns

High concreteness is often interpreted as a measure of psychological distance [13] and tweets using first-person plural (mean=4.34) tended to be marginally weaker in concreteness than tweets in first-person singular (mean=4.35) or in baseline (mean=4.38). Contributions of concreteness that vary with pronoun were negligible in the community as a whole and more obvious in the analysis of the supertweeter subcommunity described below.
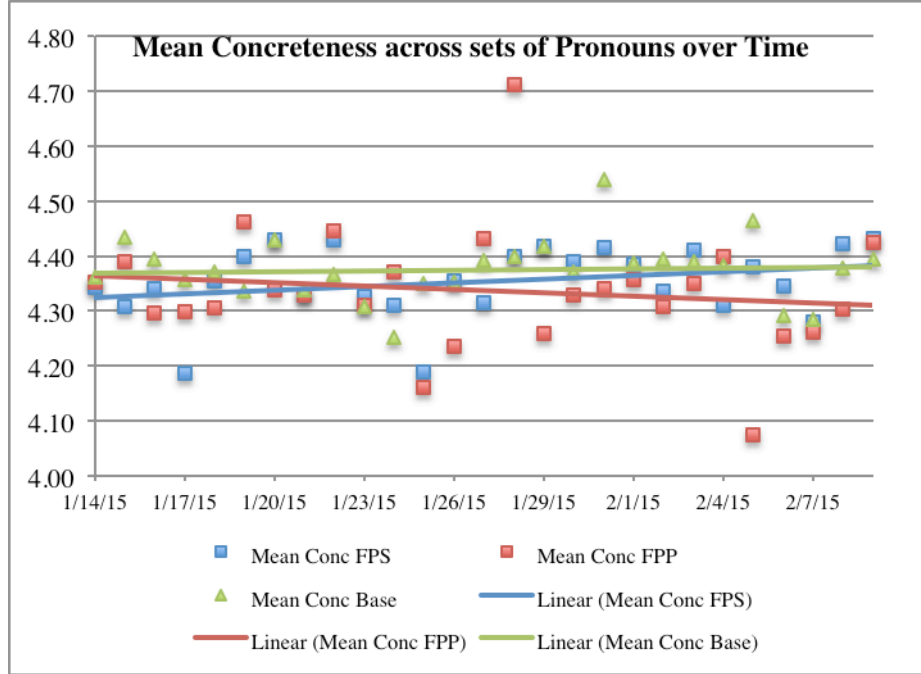
Fig. 6. Scatter diagram with regression lines over mean concreteness of words that co-occur with first-person singular (FPS) and first-person plural (FFP) pronouns and baseline (no first person pronouns).

4.5 Differences in Language Use of Supertweeters vs. the Rest

Prior research suggests that supertweeters can be likened to those at the epicenter and they tend to have higher arousal than those whose activity makes them more peripheral to the focal [4]. In order to verify such findings and others, we analyzed the differences in language use of the top 2% (by mean rate of tweeting) of the twitter users in our collection. We divided the entire collection of tweets into three separate phases of roughly equal size (in terms of number of tweets) across time. We then analyzed the lexical variation in each phase by the use of pronouns similar to our prior analyses.

Valence and Arousal Analysis (Figure 7 and 8): Super tweeters did differ from non-supertweeters with respect to the emotional content of their tweets. T-tests showed more negative tweets (lower valence) in non-supertweeters than supertweeters (2.66 vs. 2.78) but no difference in arousal (5.97 vs. 6.00).

Concreteness Analysis: Super tweeters did not differ from non-supertweeters in their use of more concrete words or fewer different concrete key words overall. These are words whose concreteness rating is in the extreme 25% of the concreteness database [11]. However, among non-supertweeters, messages containing first-person singular (4.34) or first-person plural (4.29) tend to be significantly lower in concreteness than messages without pronouns (4.38). In addition, messages using first-person plural pronouns had significantly lower average concreteness than messages in first-person singular ($p <.005$). Among supertweeters none of these contrasts for concrete-

ness were significant. Results suggest greatest psychological distance for tweets with no pronoun as well as greater distance with first-person plural than with first-person singular.
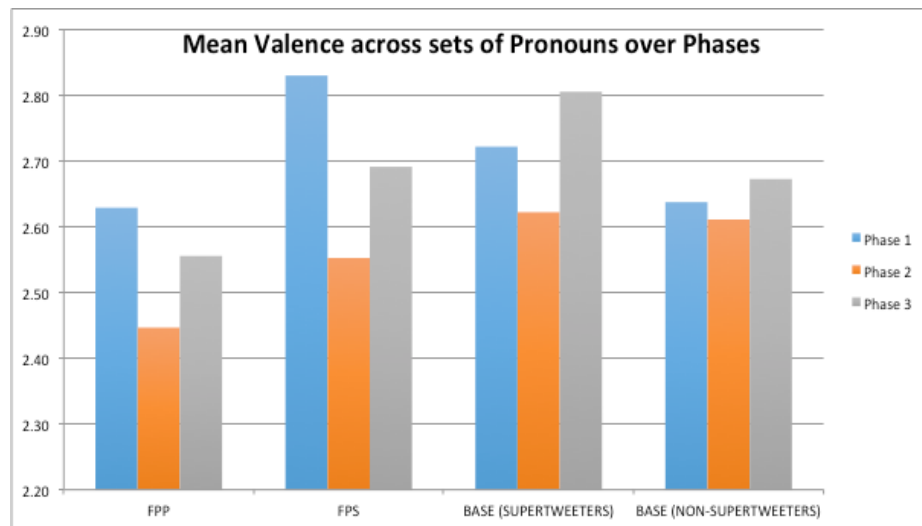

Fig. 7. Analysis of valenced language use by supertweeters across three phases and pronouns (first-person singular (FPS), first-person plural (FPP) and baseline) in our collection. The rightmost sets of bars in this chart compare these measures against the baseline for non-supertweeters.

Fig. 8. Analysis of arousal language use by supertweeters across three phases and pronouns (first-person singular (FPS), first-person plural (FPP) and baseline) in our collection. The rightmost sets of bars in this chart compare these measures against the baseline for non-supertweeters.
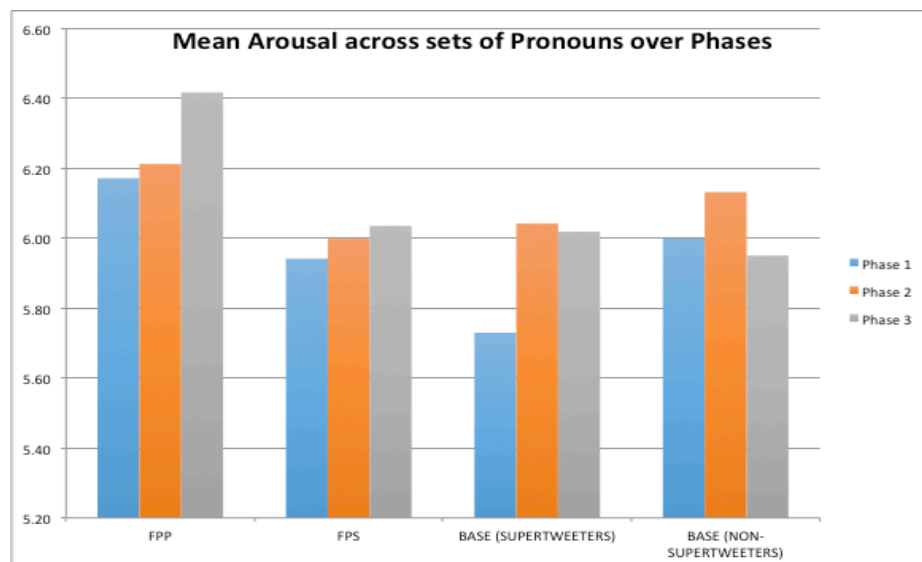
4.6   Differences in Language Use of Supertweets vs. the Rest

To learn about which tweets circulate most, we compare the composition of tweets based on frequency of the tweet message. We classify tweets based on those that are (re) tweeted most often (top 2%) and the remainder. We define tweets that circulate the most (top 2% with a retweet rate higher than 3 stdev above the mean rate of tweeting) as supertweets (N=173) and contrast their composition with that of the remainder of tweets (N=55954). Although the total number of tweets in our collection is over 400,000, we found that there are 56127 total unique tweets (14%).

In line with many studies, first person pronoun usage appears to increase in the subsequent period of events associated with threat and disruption [6]. In this case, it is first person singular pronouns that increase most. Moreover, the presence of individual differences in tweeting style suggests implicit and collective leadership roles within an online community. Most interesting was supertweets containing first-person plural pronouns were significantly lower in valence (2.55) and higher in arousal (6.38) than the baseline ($p<.05$). First-person singular supertweets did not differ from the baseline (2.69, 6.02). Like the analysis on all tweets, arousal for supertweets was significantly higher for first-person plural than for first-person singular but the difference in valence missed significance.

We examined supertweets (2% most frequent messages) separately in each of three phases (thirds) of the data. With respect to supertweets, in each case, they were more likely to contain first-person singular than first-person plural pronouns. Moreover, the asymmetry was higher in the last phase. With respect to changes in the frequency of key words in super tweets, "truth" appeared in 3% of tweets in the final phase but was negligible previously, whereas "insult/offense" and "hypocrisy" were close to 5% at one of the two earlier phases but then dropped to 1% or lower. In future work, we aim to analyze the composition of the supertweets in more detail to ascertain patterns that distinguish them from the remainder of the set.

5   Summary and Future Work

In this article, we focus on the linguistic characteristics of behavior within a network based on the diffusion of affect and other information content within tweets. We demonstrate that the tweets produced by a massive-scale network and associated with a dreadful event can be analyzed with psychological indicators such as valence, arousal and concreteness. Most consistent was that tweets with variants of the first-person plural pronoun included more high arousal words than other tweets and that arousal changed little over the course of the event. Overall, tweets with pronouns are more negatively valenced than those without and they got slightly more negative over time.

With appropriate data binning into phases, the valence of supertweeters also got more negative over time for tweets in either first-person singular or first-person plural. By contrast, in the subcommunity of supertweeters (14%) communication without first person pronouns resembled that of nonsupertweeters, in that it got less negative over time. The valence finding captures different early affective reactions towards a dramatic event in our network. We interpret progressively more negative valence in

the posts of supertweeters using first-person plural and first-person singular in conjunction with decreased negativity in the nonsupertweeters as evidence of a backlash. The increase in supertweets about truth and the decrease in supertweets about insult/offense and hypocrisy may be capturing the same phenomenon.

In summary, the relatively time constrained interactions based on individual posts could allow us to track the emergence of a sub-community that is diverging in an asynchronous manner from the community as a whole. The reversed change in the valence marks a disruption within the network that is triggered by a sub community of supertweeters whose style of tweeting deviates from that of the majority. Arousal is high throughout although patterns of lexical variation did not differ significantly be-tween the super-tweeters and the non-supertweeters. Unlike arousal and concreteness, in our data it seems that valence better captures the strength of the disruption intro-duced by a sub community on the network's collective behavior. Consistent with some previous studies, we have demonstrated how quantitative analysis of sentiment in Twitter text can be immensely helpful to uncover hidden patterns and underlying processes that may provide new insights into collective behavior within a massive-scale network.

References

1. Aragon, C. R., Chen, N-C., Feldman. L B., Kroll, J. F. Emoticon and Text Production in First and Second Languages in Informal Text Communication," In W.G. Kennedy, N. Agarwal, and S.J. Yang (Eds.): SBP 2014, LNCS 8393, pp. 217-222 (2014)
2. Brysbaert, M., Warriner, A.B., & Kuperman, V. Concreteness ratings for 40 thousand generally known English word lemmas. Behavior Research Methods, 46, 904-911. (2014)
3. Barberá John T. Jost, Jonathan Nagler, Joshua Tucker, and Richard Bonneau. Tweeting from Left to Right: Is Online Political Communication More Than an Echo Chamber? Psychological Science, 26 (10), 1531-1542. (2015a)
4. Barberá P, Wang N, Bonneau R, Jost JT, Nagler J, Tucker J. The Critical Periphery in the Growth of Social Protests. PLoS ONE 10(11): e0143611. doi:10.1371/journal.pone.0143611 (2015b)
5. Burns, A., and Stieglitz, S. Towards more systematic Twitter analysis: metrics for tweeting activities. International Journal of Social Research Methodology, 16, 91-108. (2013)
6. Chung, C. K., & Pennebaker, J. W. The psychological functions of function words. In K. Fiedler (Ed.), Social Communication (pp. 343–359). New York, NY: Psychology Press (2007)
7. Dubois, E. and Gaffney, D. The Multiple Facets of Influence: Identifying Political Influentials and Opinion Leaders on Twitter. The American behavioral scientist. 58, 10 (Sep. 2014), 1260–1277 (2014)
8. Gustafsson SendØn, M., Lindholm, T., & Sikstrom, S. Biases in News Media as Reflected by Personal Pronouns in Evaluative Contexts. Social Psychology 2014; Vol. 45(2):103–111 (2013)
9. Giglietto, F., & Lee, Y. To be or not to be Charlie: Twitter hashtags as a discourse and couter-doscourse in the aftermath of the 2015 Charlie Hebdo shooting in France. 5th Workshop on making sense of Microposts. (2015)
10. Stefanone, M.A., Lackaff, D., & Rosen, D. Contingencies of self-worth and socialnet-working-site behavior. Cyberpsychology, Behavior and Social Networking, 14, 41–49. (2011)

11. Warriner, A.B., Kuperman, V., & Brysbaert, M. Norms of valence, arousal, and dominance for 13,915 English lemmas. Behavior Research Methods, 45, 1191-1207. (2013)
12. Xu, W.W. et al. 2014. Predicting Opinion Leaders in Twitter Activism Networks: The Case of the Wisconsin Recall Election. The American behavioural scientist. (2014)
13. Trope, Y., & Liberman, N. Construal-level theory of psychological distance. Psychological Review, 117, 440–463. (2010)