Investigating Conversation Connectivity in Politically- and Socially-Charged Twitter Discussions

Devney Hamilton
devney@stanford.edu

Omosola Odetunde
omosolao@stanford.edu

Kelsey Young
christky@stanford.edu

December 10, 2014

Abstract

Twitter has frequently served as a platform for political and social discussion, playing a notable role in organizing protests, serving as a form of civil disobedience, and being used as a means of expressing outrage over social events. We seek to further research aspects of Twitter conversations (specifically their use of Twitter hashtags), which allow them to remain connected and bring together individuals and groups with differing opinions. Our hope is that, by finding these characteristics of Twitter hashtags, we can further elucidate methods of sustaining debate and discussion around political and social movements as well as bridging various communities of thought. We focus specifically on the Twitter conversations related to the August 2014 shooting of 18-year old Michael Brown, which have been centered on the hashtag #ferguson. We generate both undirected graphs and multigraphs to represent this Twitter conversation using co-occurring hashtags. We run two experiments on our graphs in order to investigate the effect on connectivity of node influence, edge presence and frequency, as well as node presence and frequency.

1 Introduction and Motivation

Though Twitter began as a platform for sharing "a short burst of inconsequential information," the site quickly made a name for itself as a medium for generating and participating in conversations of political and social natures on a global scale. The platform has been used to organize protests (e.g. the 2011 Egyptian Revolution), as a form of civil disobedience, and as a means of expressing outrage over social events.

However, Twitter has often been criticized for the transient nature of its controversial and political conversations. Large groups rally behind social and political issues, but powerful conversations can very quickly die out. We are interested in exploring the ways in which Twitter conversations around political and social justice movements can be sustained and break through fragmented communities or silos of thought.

Specifically, we are analyzing the conversation surrounding the August 2014 shooting of 18-year old Michael Brown in Ferguson, Missouri. The event quickly charged up racial and political conversations across the Twittersphere. As we have seen, US political conversations online are highly fragmented with little information passing between clusters of opinion. The #ferguson hashtag has been used for news transmission, but also as a part of highly contentious debates about racism, police violence, and other issues in the United States. As such, it is a hashtag that has the potential to cross lines of opinion and perspective. The Ferguson Twitter conversation brings up interesting questions about the connectedness of conversations, the means in which conversations can spread across disparate and diverse demographic groups as well as topic areas, and the persistence of controversial topics.

2 Problem Definition

Beginning with a database of 1.3 million tweets surrounding the #ferguson hashtag, we examine each tweet and extract the hashtags from within the text. We use the term "co-occurrence" to refer to when a pair of hashtags both appear in a given tweet. While processing our data, we also keep track of the number of times each co-occurrence, defined by a pair of hashtags, occurs. By rendering a graph of hashtags as nodes, linked with a weighted (or multi-graph) edge representing the rate of co-occurrence, we are able to look for further insight into the following questions:

1. Does the ‘strength of weak ties’ concept apply for Twitter hashtags? In other words, are tweets sharing less frequently co-occurring, or more frequently co-occurring, hashtags keeping widespread conversations like #ferguson connected?

2. What role do the most-popular hashtags, such as #ferguson, play in keeping a conversation going? In other words, if it were not for these extremely common hashtags ‘anchoring’ a conversation, would the related tweets no longer be connected, or might they be connected with a more diverse collection of hashtags?

We want to consider the role that hashtags play in keeping a conversation connected and widespread, rather than a single hashtag’s lifetime in the network. These questions connect two areas of research: the first examines what makes hashtags or blog posts viral, persistent, influential, and sticky (i.e. they catch on quickly after relatively few impressions). The second is about the global efficiency principle (the weight of a link should correlate with the number of shortest paths passing through it), the strength of weak ties , and the (alternative) dyadic hypothesis (the strength of a link is independent of its network context) [4]. Prior work on this data set has indicated that there are clear clusters or sub-conversations linked to the #ferguson hashtag [6].

3 Relevant Prior Work

Myers et al. algorithmically modeled information discussion within social network platforms, such as Twitter [3]. The researchers accomplished this by coming up with a probabilistic model to determine a node’s chances of becoming ‘infected’ (adopting the epidemic) based on its external and internal exposures.

Romero et al. also investigated information diffusion through Twitter, but focused primarily on persistence (the relative extent to which repeated exposures to a hashtag continue to have marginal effects) and stickiness (the probability of adoption based on one or more exposures) [5]. The results highlighted that political posts on average were the most persistent of all the categories while other topics such as ”games”,”sports”, and ”celebrities” would be more prominent in a random sampling, but had significantly less persistence scores.

The metrics of the phenomena from Romero et al.’s research are closely related to Agarwal et al.’s research as they looked a the effect of influence on information diffusion. Argarwal et al developed an algorithm that takes in blog information such as length of the text, number of comments, number of outlinks referenced in the post, and number of inlinks referenced to the post and outputs an ”Influence Score”[1]. This is relevant given the nature of blog sites where there is an array of features that control the stickiness and persistence of posts. In our experiment we can learn from this model various ways to curate the persistence and stickiness of different hashtags.

Cunha et al. also found that hashtag usage in tweets display preferential attachment, meaning that already popular hashtags continue to gain in popularity (“the rich getting richer” model). While they focused on tweets’ links to hashtags, we focus on links among hashtags themselves [2]. Throughout our experiment, if we observe that that hashtags with high influence scores are clustered then that shows evidence of preferential attachment. If we see a hashtag or set of hashtags re-occurring either by themselves or with other less common hashtags, then we can designate this model as being preferential. These results could relate to a trend or epidemic that propels the popularity of different hashtags. Further, using this model can benefit us by providing insight into tracking the popularity of certain hashtags and analyzing the metrics which make those hashtags popular.

The purpose of this project is to analyze the nature of influential hashtags. By sculpting a model of how to make a hashtag influential we hope to provide infrastructure for using social media to cure societal issues. For instance, Emma Pierson illustrates how disconnected clusters of people can be within the social network. Pierson displayed how Republican and Democratic activists are completely separated in terms of their Twitter society. There are so many strong ties within the clusters which means that one’s twitter feed likely only consists of viewpoints similar to oneself. Further the tweets from those users are concerns that differ from users of different political beliefs. While those concerns are posted within the cluster they do not reach people in other clusters. Pierson relates with this in stating, "So, when it comes to Ferguson, two groups with very different political and racial backgrounds ignore each other." It is likely that ones concerns aren't being resolved because they never reach the destinations that they need to. In hindsight hashtags in a tweet serve as a proxy for getting the users of differing political parties to interact. If our group can devise a heuristic for producing viral hashtags that are invariant to clustering then it may help solve this issue.

4 Data Source

We use a set of 1.3 million Ferguson-related tweets collected by researcher Emma Pierson. Just a few weeks after the Ferguson shooting, Pierson began to collect these Ferguson-related tweets and their metadata [6]. She published this set in the form of a publicly-available csv file. The tweets were collected starting 8/15/2014 and ending 9/26/2014 via Twitter's API, set to monitor the hashtags and tokens "#ferguson, #iftheygunnedmedown, #dontshoot, #handsupdontshoot, #handsup." The sampling rate was 10%.

5 Model

We rendered our data as a network with the underlying format (hashtag1, hashtag2, co-occurrence-count), where each hashtag is a node. Therefore, each node id is a hashtag id, and the edge weight is the co-occurrence-count). Given this underlying representation, we represent the co-occurrence network in the following way:

Undirected Weighted Graph Network: An edge between two nodes represents that there is least one co-occurrence of the two hashtags represented by the two nodes, and the edge has an attribute that is the count of tweets those hashtags co-occurred in. Hashtags in the dataset that never co-occurred with another hashtag are represented as zero degree nodes. We also considered a version that excluded zero-degree nodes. We did not find significant differences between these two in terms of global attributes on the graph, except, of course the number of zero-degree nodes, so we used the representation including these zero-degree nodes since we are studying the role hashtags play in connectivity on their own and in co-occurring pairs.

In order to understand the role co-occurring hashtags and individual hashtags play in the graph's connectivity, we conducted two different experiments to measure connectivity as a function of particular removed edges and nodes:

1. Experiment 1 (edge removal): We removed edges in increasing order of their weight (which is their relative co-occurrence rate). We defined a timestep to be the removal of 1/100th of edges in the original graph. At each timestep, we measured the estimated diameter, number of connected components, and the fraction of the largest connected component of the remaining graph. We repeated this process two additional times. In the second occurrence of this experiment, we removed edges in decreasing order of their weight and in the final occurrence, we removed edges in random order.

2. Experiment 2 (node removal): We removed nodes in increasing order of their frequency in the original dataset. We defined a timestep to be the removal of 1/100th of nodes in the original graph. At each timestep, we measured the estimated diameter, number of connected components, and the fraction of the largest connected component of the remaining graph. We repeated this process two additional times. In the second occurrence of this experiment, we removed nodes in decreasing order of their frequency and in the final occurrence, we removed nodes in random order.

5.1 Initial Graph Statistics

After rendering our data as a network, we performed initial statistics collection in order to gain insights that might guide us for the duration of our research. We report these preliminary statistics in Figures 1 and 2. As can be seen from these charts, our data contained a substantial number of co-occurring tweets, which made our proposed experiments a promising avenue for further exploration. Additionally, we see a very strongly connected graph with almost 97% of the undirected graph excluding zero degree nodes to be connected in a single SCC.

Since we desired to explore the connectivity of our graphs, our initial statistics confirmed that we could expect to discover interesting patterns in our dataset.

We further explore some of these statistics and their meanings in our Results and Findings section.

Number of nodes	30513
Number of edges	131174
90% effective diameter	2.816220
Number of closed triads	517995
Fraction of closed triads	0.001442
Maximum SCC size	96.8702%
Approximate full diameter	8

Table 1: Undirected Graph - Excluding Zero Degree Nodes

Number of nodes	34058
Number of edges	131174
Number of Zero Degree Nodes	3545
90% effective diameter	2.816220
Number of closed triads	517995
Fraction of closed triads	0.001442
Connected Component Size:	86.7872%
Maximum SCC size:	86.7872%
Approximate full diameter	8

Table 2: Undirected Graph - Including Zero Degree Nodes

5.2 Influence Scores for Measuring Hashtag/Co-Occurrence Rank

In addition to measuring decomposition of the graph, we wanted to keep track of the actual types of data being removed at each timestep. This was done to balance the correlation between experimental data and useful real-world application. For instance, in running the node analysis experiment and removing higher frequency hashtags we see a dramatic decay in the graph early on (Figure 3, Green Line). From first glance one might have an intuition that the graph is breaking down rapidly because the most commonly reoccurring hashtags are removed. This is not the case; it happens as result of the most "influential" nodes being removed which will be made clear within this section. First, let's make the distinction that the highest weighted (most frequently occurring) hashtag or co-occurrence does not imply that the hashtag or co-occurrence is the most influential.

Assume we have a cluster of N Twitter users and the most heavily weighed hashtag or co-occurrence appears N times, but each time it only appears to the same 1 person (without loss of generality assume N >> 1). In the same network let's say we have a hashtag or co-occurrence that appears once but is seen by N users. The most influential data point is the one which appeared to N users according to the preferential attachment theory from Cunha et al. (referenced in section 3) and the influence maximization algorithm from Agarwal et al [1][2].

The notion of an Influence Score came from Agarwal et al in their attempt to quantify the influence of a blog post. The algorithm took an additive approach by allowing for multiple parameters to influence an influence score, such as length of a post, number of inlinks, number of outlinks, and the 'goodness' of the post.[1] In quantifying the influence of hashtags and co-occurrences, we will take a similar approach. In our case let's define the Influence Score, I, of a hashtag or co-occurrence, a, as:

$$I(a) = \frac{1}{n}\sum_{t=1}^{n}\frac{\alpha_t}{\beta_t}$$

where t = single tweet that includes a, n = total number of tweets that include a, α_t = number of retweets for tweet t, and β_t = number of followers for the profile that produced tweet t . Also note that $\frac{\alpha_t}{\beta_t}$ is a *partial influence score*. The influence score is the average of all of the partial influence scores. The ratio $\frac{\alpha}{\beta}$ helps to keep the score standard over varying sample spaces. For instance if hashtag a appeared in a tweet which had 40 retweets and was shared out to 100 users, it should have the same score as if the hashtag appeared in a tweet with 4 retweets but only reached 10 users. The idea is to measure the impact of the hashtag given the number of followers. The same courtesy extends to co-occurrences. One edge case is if β = 0 and $\alpha >$ zero then we set the influence score equal to $\alpha + 1$. The logic behind this goes as follows:

Say we have N retweets and 1 follower (assuming the follower retweets once and N >> 1). The partial influence score would be $\frac{N}{1}$. Now say we still have N retweets but 0 followers then $\frac{N}{0}$ = N + 1. We say the second partial influence score here is higher because theoretically the number of retweets with respect to the number of followers is higher (which is the quantity we want to measure).

We take a partial influence score for each unique tweet a hashtag and co-occurrence appears in. To get the total influence score for a hashtag or co-occurrence, one must compute the average of its partial influence scores. Another subtlety to note is that for a unique tweet the same partial influence score is assigned to every individual hashtag and any combination of multiple hashtags (co-occurrences) in that tweet. Thus, the hashtags or co-occurrences that are most influential are the ones with the highest average influence scores. Note that this score does not depend strictly on the frequency of a hashtag or co-occurrence.

6 Results and Findings

6.1 Graph attributes

We will further expand upon the significance of various statistics within our graph (focusing on the graph including zero degree nodes) which give context to our later results and analysis. The 90% effective diameter was 2.7, the degree distribution appears to follow a power law (see Figure 6.1), and the fraction of triads that are closed is very small: .1%. These, combined with trends found in our experiments described below, suggest that the co-occurrence graph resembles a preferential-attachment model, where for any given hashtag, a tweeter is most likely to pair it with a hashtag that is already paired with many others. The intuition here is that if a tweeter wants to maximize the influence of his or her tweet, he or she will include a hashtag that many tweeters pay attention to, and more frequent hashtags are more likely to have many tweeters' attention. Entries with high influence scores are depicted in the next section about experiment findings.

6.2 Experiment 1: Edge Removal

Removing edges in decreasing order of weight produced a connectivity pattern almost identical to random removal. Removing connections between pairs of hashtags with high co-occurrence before those with low co-occurrence did not change the connectivity of the graph.

Removing edges in increasing order of co-occurrence weight produced a pattern different than that of removing nodes randomly or in decreasing order. This may suggest that relatively infrequently co-occurring hashtags may play a partial role in graph connectivity. Removing connections between pairs of hashtags with low co-occurrence before those connections with higher co-occurrence caused an acceleration in the disintegration of the graph about halfway through the decay process. We randomly sampled 10 pairs of hashtags whose link was removed at each timestep. Examples of pairs whose link was removed early were relatively peripheral to the discussion: (autism dontblameautism), (drones, fisa), (fluiddymamics, yikes), (assaultrifle, jihad), (icebucketchallenge, ukraine), (ferguson, ramproud), (comcast, nsa), (darfur, thailand).

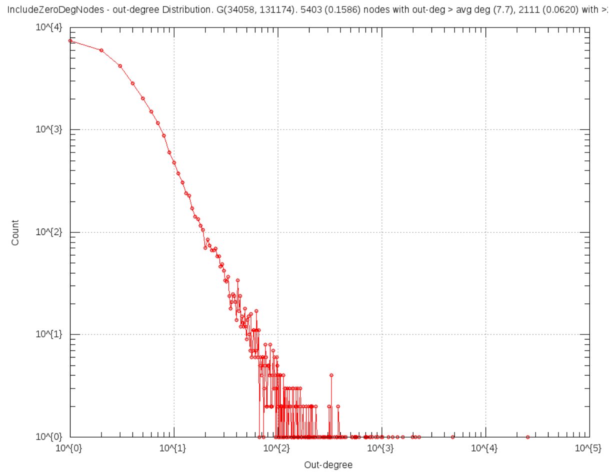

Figure 1: The degree distribution of the co-occurrence graph appears to follow a power-law.

Those removed in the middle, causing the accelerated decay in the 'increasing weight' removal order, either relate #ferguson to other discussions such as (kiev, ferguson), (drones, revolution), (womensequalityday, michaelbrown), (mikebrown,neonazizionst); or represent less popular hashtags within the #ferguson conversation, such as the pair "ripmikebrown, stopdontshoot" which is similar to the more popular "mikebrown, handsupdontshoot" tag pairing, or pairs including lesser-known names of victims of police violence like Kajieme Powell, John Crawford, and Ezell Ford. Tag pairs removed at the very end are pairs of well-known hashtags used to refer to the movement against police violence in black communities like (blacklivesmatter, every28hours) and (handsupdontshoot, ineedanswers). This suggests that at some scale, relatively rare pairs of hashtags are connecting relatively unrelated conversations, like "kiev," "womensequalityday," and "neonazizionst" that would otherwise never be connected. However, it is not a strong effect, so there are clearly other factors at play in graph connectivity, suggesting that the strength of weak ties hypothesis applies to hashtag co-occurrence much less than it applies to human communication networks.

6.3 Experiment 2: Node Removal

Removal in decreasing order of frequency:
Removing nodes in decreasing order of frequency (in green), meant that the most common hashtags such as #ferguson are removed first. The behavior here suggests that hashtag co-occurrence follows a preferential-attachment pattern, with hub nodes (rather than triangles) keeping the graph connected. By examining the green lines in Figure 6.3, we can see the graph breaking into smaller, disconnected subgraphs much more quickly than with either increasing-frequency or random note removal, and edge-removal strategies. This suggests that the most frequent hashtags keep the original graph connected, and actually act as hub nodes. With decreasing removal order, the size of the maximum connected component indicates that the graph is broken into very small components after about 12 timesteps (green line in 3rd plot). In the diameter plot, we get another perspective on how the graph rapidly falls apart with the decreasing removal pattern: as the size of the maximum connected component falls steadily in the first 10 timesteps, the diameter steadily increases before crashing around the time all connected components become very small. The number of connected components rises in the beginning of the process, as the discussion breaks into sub-discussions, and then falls as lower-frequency nodes that form small connected components are removed; high frequency hashtags are of course more likely to co-occur and therefore be in larger connected components. Low-frequency nodes are more likely to start off in small connected components (including those of size 1). These measurements indicate that a small number of relatively significant sub-discussions become disconnected from each other with the removal of the most frequent hashtags. Those sub-discussions seem to be internally anchored by

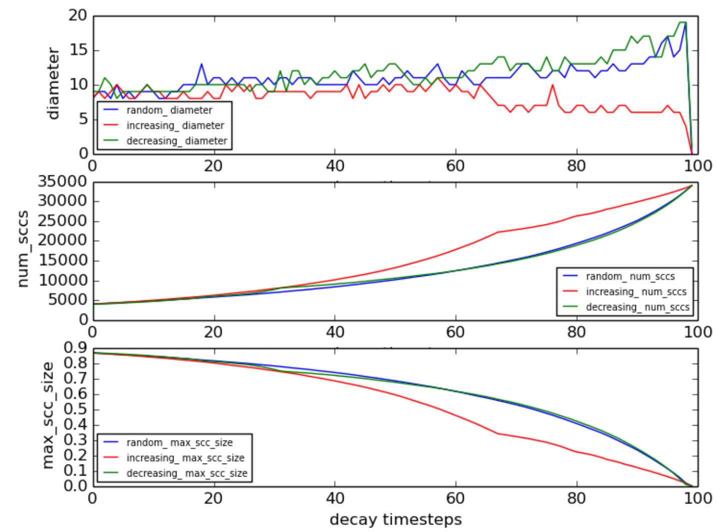

Figure 2: This plot displays the measurements of diameter, number of connected components (SCCs), and the size of the largest connected component relative to the graph at that time, taken at each timestep in edge removal. The red line represents removal in order of increasing weight, the green decreasing, and the blue random.

'minor hubs' with mid-level frequencies, and then after those are removed the diameter drops as well because the remaining nodes are in very small connected components. This experiment strongly suggests a preferential-attachment structure with a few classes of 'hubs,' consistent with the lack of closed triads and power-law distribution of degrees.

Removal in increasing order of frequency:
Removing nodes in this order meant that the size of the largest connected component only increased, and the number of connected components only decreased. This suggests a high correlation between a node's frequency in the original dataset of tweets and its centrality in the co-occurrence graph. Removing in this order was equivalent to repeatedly removing nodes with the lowest centrality at each timestep. The fact that the diameter remained fairly constant until the very end of the decay process is consistent with this pattern.

Removal in random order of frequency:
The steady diameter, similar to the increasing order node removal strategy, suggests that the structure of the graph is similar at many scales; except where the drop off in max-scc size indicates the graph quickly 'fell apart' at the loss of a key hub node (likely #ferguson).

7 Discussion

Given our findings from these experiments, we will look into some of the initial motivating questions we had for our research.

1. Does the 'strength of weak ties' concept apply for Twitter hashtags? In other words, are tweets sharing less frequent co-occurring hashtags keeping widespread conversations like #ferguson connected?

   The 'strength of weak ties' concept was not definitively displayed in our experiments. It seems that

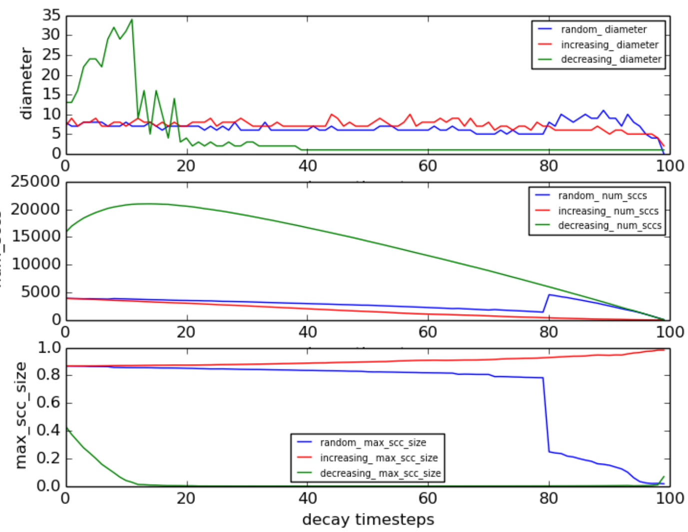

Figure 3: This plot displays the measurements of diameter, number of connected components (SCCs), and the size of the largest connected component relative to the graph at that time, taken at each timestep in node removal. The red line represents removal in order of increasing frequency, the green decreasing, and the blue random.

although prior on, we found clearly distinguished communities in the co-occurrence of hashtags of our dataset, suggesting 'sub-discussions' [6], these clusters are not like the communities in the cellphone graph described by Onnela et al[4]. The cellphone graph involved more 'small-world' effects in local clusters, with triangles. Furthermore, during the course of our research project, Pierson published further results from her datasets which showed that the #ferguson discussion is clearly split into two large clusters with very little connection between them [7]. Rather than displaying the strength of weak ties in human social networks, it is more likely that the hashtag co-occurrence graph resembles a preferential-attachment pattern, with hubs for each of the very large clusters.

2. What role do the most-popular hashtags, such as #ferguson, play in keeping a conversation going? In other words, if it were not for these extremely common hashtags 'anchoring' a conversation, would the related tweets no longer be connected, or might they be connected with a more diverse collection of hashtags?

   Our experiments displayed results similar to those of Agarwal et al previously discussed. Calculating an influence score for each individual hashtag and/or co-occurrence helped to pinpoint some key features in the "hub" nature of a hashtag. #ferguson held the highest influence score because it was both most widely viewed as well as most often tweeted. This seems reasonable since #ferguson seems to very directly reference the issue at hand. It is also quite short and seemingly neutral. #ferguson is not inherently pertaining to Republican or Democratic or Conservative or Liberal views on the subject, therefore it is a medium through which users from both ends of the spectrum can interact and discuss the same event. In our experiments, we found that highly influential nodes, such as #ferguson, serve as hubs in the network and their removal very quickly disconnects the graph.

7.1 Implications

Two of our findings support Pierson's suggestion that users of the #ferguson hashtag are communicating little with those of differing perspectives. The large proportion of 'open triads' in the co-occurrence graph means that for any two hashtags (e.g. "tcot" for "top conservatives on twitter" and "tlot" for "top liberals on twitter") that both co-occur with the #ferguson hashtag, it is very unlikely those two every co-occur. Therefore it's unlikely that tweeters who use #ferguson will see tweets from other 'sub-discussions' within the #ferguson conversation. The second finding that supports this is that removing high-frequency nodes disconnected the graph very quickly, as opposed to removing edges with weight of co-occurrence.

Overall, it seems that the #ferguson hashtag does cross lines of perspective, but doesn't necessarily carry differing ideas with it. It acts as a hub, and perhaps 'hub' hashtags do sustain conversations because as their co-occurrence with other hashtags increases, so does their likelihood of co-occurring with additional hashtags. Once these 'hub' hashtag are removed, the Twitter conversations (and represented network) quickly become fragmented, often representing silos of discussion and thought. As we displayed in our experiments, these types of hashtags have the ability to keep disparate conversations connected.

8 Future Work

Further research might involve closer investigation into preferential attachment over time. We can investigate if preferential attachment leads to scale-free conversation structures.

In the next stages of this research, it would be interesting to explore the time effects of co-occurrence in our graph. Since our data was specifically collected starting on a particular day, we do not have fully reliable data about the presence and co-occurrence frequency of hashtags over time. It is possible that certain hashtags were frequently co-occurring before the collection period or in general serve as common co-occurrence pairs even beyond the Ferguson conversation. Our dataset would not capture these information.

In future work, we might want to start collecting data continuously in order to capture the Tweets immediately as the event occurs. If we did this, we could then answer some additionally interesting questions, such as those related to the decay factor (or persistence) of "hub" hashtags. How quickly do they decline? How does their decline mirror the decline of other hashtags? How quickly is the overall network disconnected during their decline? Given our previous findings that hub hashtags anchor a conversation and hold together fragmented views, it is particularly interesting and important to determine how persistent these binding hashtags can be.

Additionally, we were not able to track the number of exposures a given Twitter user had to a particular hashtag. If we had such information, we could explore the "stickiness" of various hashtags (i.e. how many exposures a user must have before adopting a given hashtag). This could provide interesting information about how quickly hashtags can spread, limiting factors in their spread, and a prediction of which hashtags we expect to more readily co-occur together given their individual "stickiness."

Since we are interested in the connectedness and sustainability of political and social conversations in Twitter, exploring persistence and stickiness further would be great steps for future development and insights.

References

[1] Nitin Agarwal , Huan Liu , Lei Tang , Philip S. Yu, Identifying the influential bloggers in a community, Proceedings of the 2008 International Conference on Web Search and Data Mining, February 11-12, 2008, Palo Alto, California, USA

[2] Evandro Cunha, Gabrial Magno, Giovanni Comarela, Virgilio Almeida, Marcos Andre Goncalves, Fabricio Benevenuto. Analyzing the dynamic evolution of hashtags on Twitter: a language-based approach. ACM LSM '11 Proceedings of the Workshop on Languages in Social Media. 2011.

[3] Seth A. Myers , Chenguang Zhu , Jure Leskovec, Information diffusion and external influence in networks, Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining, August 12-16, 2012, Beijing, China

[4] J.-P. Onnela, J. Saramaki, J. Hyvonen, G. Szabo, D. Lazer, K. Kaski, J. Kertesz, A.L. Barabasi. Structure and tie strengths in mobile communication networks. PNAS, 2007

[5] Daniel M. Romero , Brendan Meeder , Jon Kleinberg, Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter, Proceedings of the 20th international conference on World wide web, March 28-April 01, 2011, Hyderabad, India

[6] Emma Pierson. How to study the rage of millions of people, Blog post, August 19, 2014.

[7] Emma Pierson. See how red tweeters and blue tweeters ignore each other on Ferguson, Blog post, 26 November 2014.