Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media

Conversational Shadows: Describing Live Media Events Using Short Messages

David A. Shamma, Lyndon Kennedy, Elizabeth F. Churchill
Yahoo! Research
aymans@acm.org, lyndonk@yahoo-inc.com, churchill@acm.org

Abstract

Microblogging concurrently with live media events is becoming commonplace. The resulting comment stream represents a parallel, social conversational reflection on the event. Although not formally 'attached' to the actual event stream itself, we demonstrate it is possible to establish a relationship between the two streams by mapping their structural properties. In this article, we examine: How do people produce and consume real-time commentary? And how does the structure of commentary and conversation change in response to moments of interest? Using a dataset of 53,712 Twitter posts, or tweets, sampled during the inauguration of Barack Obama in January 2009, we develop methods for exploring these questions. We find that short message activity reflects the structure and content of this media event. Specifically, messages directed at large audiences can serve as broadcast announcements, while variations in the level of conversation can reflect levels of interest in the media event itself. Finally, we present some implications for the design of future tools for a variety of users ranging from consumers to journalists.

Introduction

On January 20, 2009, Barack Hussein Obama became the 44th president of the United States of America as the world watched locally from Washington DC and remotely on broadcast TV and a host of available live Internet streams. News websites streamed live video to desktop browsers and video streaming services fed live streams to mobile devices. While watching, people concurrently updated their online photos, blogs, and microblogging services.

This paper focuses on the posts, or *tweets*, people made using Twitter, which constitute comments or annotations on the live media event that was the Presidential election. Our interest, for the purposes of this paper, is not the media event per se but the conversational shadows the event cast in the form of the Twitter-based comment stream. We describe a method for finding social and conversational insights about the event. We also examine how the structure of the group conversation online reflected the structure of the media event that was taking place offline. We demonstrate it is possible to infer the structural properties of an offline, *live* media event (e.g., lulls in the action, moments of focused engagement versus intermissions, etc) through the shape of the conversational flow over time in short-messaging services.

Owing to the short post size, it is easy to contribute Twitter messages from dedicated applications and web browsers on mobile devices as well as desktops. Individuals use short messaging services like Twitter to broadcast messages which convey opinion, presence, or awareness to a set of subscribers (Naaman, Boase, and Lai 2010). Several events, such as the 2009 MTV Video Music Awards, invite Twitter participation from the audience and TV viewers alike[1]. This use of micro-messages has fueled commenting in real-time—shifting the idea of media sharing and commentary from asynchronous/near-synchronous to fully synchronous, concurrent conversation around a common referent.

For example, on YouTube, media is first uploaded and then annotated in various ways (favorites, comments, tags, etc). By contrast, tweeting during live events produces a stream of comments that are disconnected from their referent; they are thus only annotations in the sense that they have a common referent, but they are in no way associated with their referent except by inference. One interesting aspect of this kind of concurrent commentary that stands in contrast to social media sites is the event to which the conversation is oriented or directed may not itself be available for view—thus for those unaware of the shared referent, the conversation is decontextualized.

First, how do people produce concurrent commentary, broadcasting and sharing short messages? How can a tweet address another person or a set of individuals in a conversational manner, rather than as a broadcast statement? This article addresses the methods we have derived to answer these questions, and demonstrates events cast a discoverable conversational shadow in online short-message contexts.

Background

Twitter has become an essential tool for gathering comments around events and news. In 2008, Current TV began displaying live tweets on the their TV channel and webcast[2]. This

[1] `http://www.mtv.com/news/articles/1621053/20090909/index.jhtml`

[2] `http://current.com/topics/88834922_hack-the-debate/`

program promoted and fueled the concurrent usage of microblogging commentary around live events. In 2009, Facebook and CNN ran a similar program during the inauguration. Many news channels routinely call for op-ed questions and commentary in the form of tweets.

To date, numerous video sharing websites allow for asynchronous commenting of videos. These comments do not refer to subparts of the video, but are comments on the video as a whole; some sites, like Hulu, offer the ability to clip a segment and share it via email or other social networks. Cesar et al. (2008) have conducted a series of studies involving media clipping, annotation and sharing. Their test subjects reported feeling closer and more connected while sharing media with each other. In a study of synchronous video watching online, Shamma and Liu (2009) found clear patterns to the conversation: people chat, then watch the video and wait for the video to finish before resuming conversation, and chatting at length. Recently, Shamma et al. (2009) also found fluctuations on the volume of Twitter traffic to reflect points of heightened interest in a televised debate; when points of general interest occurred, the Twitter commentary traffic decreased, and during what could be constituted as lulls, the volume of traffic increased.

Traditionally, in the case of Twitter, reciprocity has only been measured as a followers-to-following, relationship and not based on reciprocal engagement. In other words, reciprocity is measured as symmetrical friend relations in a social graph. When it comes to friendship reciprocity, Java et al. (2007) found Twitter users had a higher reciprocity than bloggers (Shi, Tseng, and Adamic 2007). Honeycutt and Herring (2009) examined the usage of the @mention notation in Twitter as the basis for understanding how conversational a user might be. In our work, we are addressing the problem of how to identify events taking place to a crowd via the patterns of social activity, both direct and peripheral participation (bystanders, lurkers, readers and so on). We are interested in being able to discern that something *must* be going on from the "swell" of crowd activity, but also to broadening the notion of participation to consumers and producers.

Study

For our study we posed the following questions: What happens conversationally when people post short messages during an event? Can we discover moments of interest? Can we find the moment where people were focused more on the event rather than their ongoing conversations? Can we learn about a media event's structure by looking at a large collection of concurrent and post-event posts?

From these questions, we investigate two hypotheses which stem from the human centered insights represented in the work of the previous section:

(H1) Activity from users with high follower counts are media-broadcaster announcements. Their activity is an indicator of a public event happening at that moment.

(H2) Points of low conversation will signify the start of an event or segment. Points of high conversation will signify the end of an event or segment.

To investigate these hypotheses, we analyze activity on Twitter during the 2009 inauguration of the Barack Obama.

Data Collection

For this study, we acquired access to Twitter's Data Mining Feed (now deprecated). This service allowed polling for 600 tweets per minute from the public time line. To study Twitter activity during the inauguration, we collected an hour and a half sample of tweets around President-to-be Obama's swearing in and his following speech: approximately 11:30 AM to 1:00 PM Eastern Standard Time, January 20, 2009. While, the ceremony itself took place from noon to 12:30 PM, we wished to study activity leading up to and after the event itself. Public timeline sampling in this manner brings a more so realistic snapshot of all of the social activity. We also used video and closed captioning of the swearing in and proceeding address from CSPAN's YouTube page[3].The total polled sample pulled 53,712 tweets from the public time line in the relevant hour and a half window. It is important to note that, as this provides a sample from the overall time line, not all of the tweets pulled were about the inauguration.

Roles and Features

Owing to sheer size of the datasets involved, examination of logged activity requires considerable data management and principled reduction through sampling and compression. Therefore before these data could be processed and analyzed, one must determine the structural properties of the data. Within a tweet, a user can specifically mention another user by prefacing their username with the "@" symbol. In Twitter, this @mention format serves three purposes. First, it creates a hyperlink to the other user's Twitter feed. Second, a user can keep track of their personal mentions to see who might be talking about them. Finally, the two or more users may @mention each other as a conversational mechanism to thread a set of messages. In addition to explicit referencing, a user can opt to subscribe to another user's posts. This is referred to as *following*. As a result, any user has a set of *followers* that subscribe to their feed. In our sample, 12,573 tweets (23.4%) mentioned another user. Often, users re-post, *retweet*, other users' posts. Approximately 1.93% of tweets in our sample are retweets.

We are interested in examining the entire sample of tweets, as well as, identified inauguration relevant tweets. Having an overall sample from the public time line allows for one to iteratively examine what might be relevant in the overall public time line traffic. To find the relevant subset, we identified a set of related terms and tags, terms that start with a # symbol. To do this, we examined a rank ordered list of all the words and selected the most frequent on-topic terms; this produced the following list: #inaug09, #current, barack, obama, biden, president, whitehouse, bush, dub (the latter referring to the outgoing President George W. Bush). We then created an inauguration subset of 13,370 tweets that contained at least one of these terms.

[3]http://www.youtube.com/watch?v=nNIEduEOw

Analysis

Having addressed some of the structure of tweets and their conversationality, we can now begin to investigate how these features are utilized during the inauguration. We examine the swearing in and opening address of President Obama plus the 30 minutes prior and after. In the 90 minutes, this includes the Bush and Obama families proceeding to the platform. A small performance that is followed by the swearing in of the vice-president and president. The inaugural address concludes the ceremony. The former president Bush then departs via a helicopter as Obama signs presidential papers. The timeline, along with the time-stamped video itself will serve as a baseline for localizing themes and event onsets.

This paper draws from the previously mentioned work of Shamma et al. (2009) In their study, they queried for tweets using three tags from a political debate and used the volume of tweets per minute as a mechanism for finding segments of the debate. Their approach varies from our study in two respects. First, we are interested in the overall representation of Twitter's public time line with respect to the media at hand—that is, if someone was not tweeting about the inauguration, did their tweeting practice nevertheless change? This would assume that the public traffic on Twitter was affected by Obama's inauguration, regardless of the content of the individuals' tweets or of the focus of their attention. Second, we also investigate all tweets related to the media event that are not restricted to simple #hashtag or @mention usage. This is an effort to understand usage by less proficient, or routine, users of Twitter. Essentially, the usage of the # and @ nomenclature comes from a subset of users. For example, in our dataset, 8 tweets contained the string "#barack", 71 tweets contained "@barack" and 1,157 tweets contained "barack" from a case-insensitive matching.

Followcasters Each user on twitter has a follower count which represents the number of people explicitly listening to that feed; this count is 0 when a new user signs up. The number of followers any user has can describe the role that he or she (or that organization) engages in on Twitter. While the follower count is not in the dataset of postings, these numbers can be determined easily. The distribution of users by their number of followers is log-linear. Few users have over 1,000 followers ($\mu = 242.9$, $\tilde{x} = 62$).

Examining the follower count by minute, one can estimate the subscriber population at a particular moment in time. To do this, we first aggregate all the twitter users tweets by minute. Then we compute the sum, mean, median, and maximum number of followers from each minute's users. We will investigate the maximum follower count by minute to identify dominant tweeters or followcasters who could be announcing event onsets, hypothesis (H1). That said, of the 13 users in the upper quartile tail, one user's follower count dropped by one only to increase by two followers within the 90-minute sample window. In the upper quartile tail, 19 tweets came from 13 users. Of these users, only two users (Robert Scoble and CNN Breaking News) were the outliers in the upper quartile tail of the overall distribution ($> Q_3$); Scoble, a prominent blogger in the San Francisco Bay Area has 49,485 followers and CNN Breaking News 86,631 followers. Scoble relays a quote from Current TV's news director:

> 'We are running low on Tweets,' Current TV's news director @mgphritz says, asking his team to post more tweets about Obama's speech #current

CNN Breaking News tweeted at 1:22:18 EST a quote from the address:

> Obama: 'America must play its role in ushering in a new era of peace.'

In both of these examples, the act of posting a quote serves to extend the dissemination of the quoted message, and this reinforces or underscores it. For those who many have missed the moment itself, this acts as a form of replay or catch-up. It may also perform a bookmarking technique, for the media reporters to echo during a recapitulation.

Segmentation & Onset Followcasters represent a small portion of our sample; many users ($\tilde{x} = 62$, $IQR = 169$) in our sample had fewer than 100 followers. Additionally 23.4% of the posters conversationally @mentioned another user. To investigate our second hypothesis (H2), we will examine how the conversation fluctuates during the course of the media event. H2's intuition comes from previous research which found people to be more conversational towards the end of a video (Shamma and Liu 2009). The inauguration was an all day continuous event. We are focusing on the actual swearing in and following speech. Can we identify the overall level of interest on a particular person or topic from the captured conversation?

We begin by mapping moments of low conversation to points of event onset during the inauguration and moments of high conversation to that segment's ending. In other words, we investigate if people will post less conversation content at significant moments and post more conversation content at the end of a segment. Periodic swells in volume should depict logical *breaks* in the event, thus reveal candidate segmentation points. Our dataset was taken from a linear rate feed; the number of messages per minute is more or less constant (≈ 596.69 tweets per minute) and not periodic. Therefore, examining overall volume by minute does not work. However, the volume of directed conversations, @mentions, varies over time ($\mu = 139.7$, $\sigma = 34.48$), see Figure 1. At 12:05 PM, only 37 tweets ($< \mu - 2\sigma$) contained an @ symbol. Given this drop is present in the entire sample of tweets, we speculate the majority of the sample was focused on the inauguration regardless of the content of their tweets.

Examining the volume of the inauguration-only tweets shows an increase at the moment of oath. Despite this volume spike, we still see a @mention drop at the start of the oath which corresponds at 12:05 PM. Inauguration tweets with an @ mention also shows a drop in conversation around the swearing in of the Vice President. Figure 1 shows each groups respective volume. Additionally the overall volume of @ mentions tweets increase over time.

The average number of characters typed per tweet by minute is 70.628 ($\sigma = 7.082$). This is computed from a subset which excludes all tweets with an @mention to keep

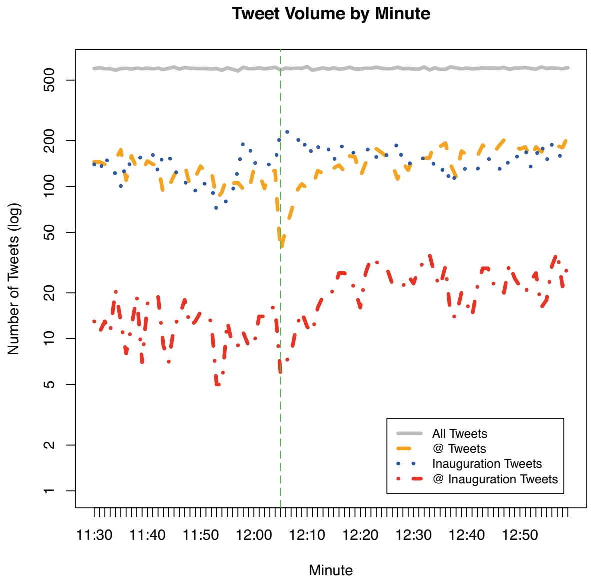

Figure 1: The number of tweets over time by minute. The overall sample averaged ≈ 596.69 Tweets per minute. *Inauguration tweets* is a subset of the dataset containing top related inauguration terms. The volume of Inauguration tweets increased during the oath of office. Of the tweets with an @ mention, there is a drop at 12:05 PM, where the swearing in of Obama began.

the username reference from artificially inflating the number of characters typed and to remove any inflation or deflation that might result from a tweet being directly conversational. From 12:04 PM to 12:05, the average number of characters per tweet drops from 76.525 to 45.697. The average number of characters typed at 12:05 p.m. is $< \mu - 3\sigma$. Similar to the drop in @conversations, we assume people were typing less as they were paying attention to the inauguration. There is a strong correlation between the average number of characters per minute and the number of @mentions per minute ($\rho_{X,Y} = 0.7919$, $p < 0.001$).

Discussion

From our hypothesis (H1), we were able to find broadcast media announcements from users with high follower counts. This examination needed to be performed on the inauguration subset of tweets in order to be relevant to the media event and to remove false positives. In the overall Twitter set, some media broadcasters, identified with a follower count $> Q_3$ in the distribution, were announcing other non-inauguration news. Examining @mention conversation to find event onset, from hypothesis H2, showed mixed results. Low conversational volume aligned to event onsets within our sample. Points of high conversational volume existed, but were not particularly aligned to the overall event time line of the inauguration. In effect, the drop in @mentions becomes a conversational gasp where one's attention is not directed to others in the network.

Future Work

We have explored two methods for discovering and understanding the conversational shadows that result from twittering in response to a live media event. This process begins with identifying several features which are conversationally salient. While we have begun to find advancements in onset detection and followcasters, we wish to identify trending and conversational topics. Additionally, we wish to investigate temporal effects on social network centrality. The threaded conversation of @user mentions could provide subtopic trends amongst small networks of people. Further, we wish to examine real-time event tracking, where the number of tweets and the size of the network would grow over the lifetime of the event and beyond. This would prove more effective in a different event genre, like a sporting match, where a plurality of actors engage in different roles over the course of the event. This combination of sub-graph centrality and real-time event tracking suggests a variety of tools, for journalists and consumers alike, to follow event news and commentary. A geographic filtration of the tweets may even provide local versus global reporting and identification of salient conversations. For example, what may have temporally sustained interest in Boston may not have any interest at all in Austin. Finally, our methods can be expanded to non-televised events. In particular, longer scale media events have no single media object for reification, such as the entire election process or an ongoing investigation. We wish to investigate how these larger scale events can be understood by examining a larger social conversation coupled with a collection of media like news articles, videos, and photos.

References

Cesar, P.; Bulterman, D. C.; Geerts, D.; Jansen, J.; Knoche, H.; and Seager, W. 2008. Enhancing social sharing of videos: fragment, annotate, enrich, and share. In *MM '08: Proceeding of the 16th ACM international conference on Multimedia*, 11–20. New York, NY, USA: ACM.

Honeycutt, C., and Herring, S. 2009. Beyond Microblogging: Conversation and Collaboration in Twitter. *Proc 42nd HICSS*.

Java, A.; Song, X.; Finin, T.; and Tseng, B. 2007. Why we twitter: understanding microblogging usage and communities. In *WebKDD/SNA-KDD '07: Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*, 56–65. New York, NY, USA: ACM.

Naaman, M.; Boase, J.; and Lai, C.-H. 2010. Is it really about me? message content in social awareness streams. In *CSCW '10: Proceedings of the ACM 2010 conference on Computer supported cooperative work*. Savannah, GA: ACM. *Forthcoming*.

Shamma, D. A., and Liu, Y. 2009. *Social Interactive Television: Immersive Shared Experiences and Perspectives*. Hershey, PA, USA: Information Science Publishing. chapter Zync with Me: Synchronized Sharing of Video through Instant Messaging, 273–288.

Shamma, D. A.; Kennedy, L.; and Churchill, E. F. 2009. Tweet the debates: Understanding community annotation of uncollected sources. In *WSM '09: Proceedings of the international workshop on Workshop on Social Media*. Beijing, China: ACM.

Shi, X.; Tseng, B.; and Adamic, L. 2007. Looking at the blogosphere topology through different lenses. *Ann Arbor* 1001:48109.