

Twitter, Information Sharing and the London Riots?

by Emma Tonkin, Heather D. Pfeiffer and Greg Tourte

Knowledge Management of Social Networks

EDITOR'S SUMMARY

The prolific commentary disseminated via Twitter on the riots in London and other British cities in August 2011 has given rise to the question of whether their reflection in such social media forums may have added to the unrest. Investigators analyzed 600,000 tweets and retweets about the riots for evidence that Twitter was used as a central organizational tool to promote illegal group action. Results indicated that irrelevant tweets died out and that Twitter users retweeted to show support for their beliefs in others' commentaries. Tweets offered by well-known and popular individuals were more likely to be retweeted. In the case of the British riots, there is little overt evidence that Twitter was used to promote illegal activities at the time, though it was useful for spreading word about subsequent events.

KEYWORDS

Twitter	computer-mediated communications
social networking	sociocultural aspects
social web	

Emma Tonkin was originally educated as a physicist, but has since diversified. She works as a researcher at UKOLN, University of Bath, United Kingdom and can be reached at e.tonkin@ukoln.ac.uk.

Heather D. Pfeiffer is college research faculty at New Mexico State University in the Klipsch School of Electrical and Computer Engineering. She is interested in knowledge engineering and management, database management and modeling of domains of information using automatic ontology building. She can be reached at hdp@cs.nmsu.edu.


Greg Tourte is a research associate at the School of Geographical Sciences at the University of Bristol, United Kingdom. He currently specializes in large climate dataset management, distribution and visualization. He also maintains a supercomputer and the associated data processing server farm in his spare time. Greg can be reached at g.j.l.tourte@bris.ac.uk.

The recent riots in London and other cities in England have inspired a great deal of interest in the role of social media in incitement, organization and analysis of events. Initial reactions to the events saw many commentators laying blame upon social networks such as Twitter, suggesting that networks are being used to incite social unrest and that their absence would reduce the likelihood of ongoing unrest.

The microblogging tool Twitter has become a popular service, widely used worldwide. Its popularity has grown rapidly in the United Kingdom, where the traditional media have recently begun to use it both as a source for information and as a dissemination platform. From its beginning, it has been the subject matter of occasional BBC news items; for example, the BBC now refer to Twitter regularly as a news source (see Figure 1). It is possible to see a cyclic element in the relationship between Twitter and the national UK media. On the one hand the traditional media are driven by various factors, including events, existing discourse and public opinion (as represented by Twitter among other sources). On the other, as the service becomes entrenched in journalistic practice, some perceive the flip side of the coin, that Twitter content is driven by traditional media and traditional journalism – thus the potential cyclic element.

As well as functioning as a forum for comments and discussion, Twitter also provides the retweet mechanism; relaying another individual's tweet through one's own account and to one's own followers is "the key mechanism

FIGURE 1. Mentions of Twitter in articles on the BBC news website



for information diffusion in Twitter” [1]. Boyd et al. [2] suggest that retweeting is part of a complex conversational ecology – more visible Twitter participants “retweet others and look to be retweeted” – and is a means of participating in a “diffuse conversation.” The frequency of repetition of a message does not imply accuracy; indeed, Boyd et al. [2] describe a case in which an inaccurate story received frequent retweets, while a published correction of the story did not.

What does it mean to retweet a message without introducing one’s own interpretation – repeating someone else’s tweet “with a straight face”? Twitter participants frequently provide a short preface to retweeted messages. A few examples of this sort of usage include:

- One or more exclamation marks, such as: ‘!’ or ‘!!’, apparently conveying shock or raised-eyebrow surprise
- “<3”, an ASCII heart symbol recalling the omnipresent “I <3” t-shirt/ bumper sticker slogan
- Brief textual comments such as “Anyone?” “Best yet:” “Couldn’t agree more”

In this way, the retweeter is able to say, “I read this – and I would like to indicate the extent to which I agree/disagree.” Under some circumstances, it seems, Twitter can operate as an echo chamber for rumor/gossip. As we will see in this article, it can also serve as a stage for humor or drama, and it can help to amplify humanitarian impulses. There is a temptation to see retweets as a “dramatic performance” [3], a sort of digital grandstanding before one’s friends and acquaintances; yet to adopt this description wholesale would be to trivialize a very real interaction.

It is useful to note the context in which the tweet is placed by the author. Terms starting with a #, or hash symbol, known as #hashtags, are treated specially by Twitter, enabling streams of tweets containing the same #hashtag to be read in sequential order by any user. This device is often used to create the effect of a topical chatroom dedicated to the subject of that #hashtag. Examples include essentially innocuous activities like games, such as #oneletteroffmovietitles (movie titles with one altered letter, radically changing their meaning) – for example, The Empire Strikes Jack, The Incredible Sulk and so forth.

New performance #hashtags appear often, and it is common for them to play on issues currently popular in traditional media. For example, the discovery that a journalist for the UK’s *Independent* newspaper had created partially fabricated interviews based on material lifted from elsewhere led to the appearance of a #hashtag, #interviewsbyhari, containing brief extracts generated using the same methodology. Any number of individuals may perform freely on this impromptu #hashtag stage. One @RichardLucas3, for example, contributed, “Shakespeare sharpened his quill, turned to me, and enquired, ‘Shall I compare thee to a summer’s day?’ #interviewsbyhari.”

Another, @betapolitics, proposed, “As Neil Armstrong came bounding over to me he said: ‘One small step for man, one giant leap for mankind’ #Harigate #interviewsByHari.” These examples demonstrate a further characteristic of #hashtags. More than one may be referenced within a single tweet. The practical limitation is the character count limit of Twitter, which penalizes excessive #hashtag use. It is not, however, uncommon for two #hashtags to be in use simultaneously for a period of time.

In general users on Twitter only see the content posted by friends – in other words, information from approved sources. Hashtags provide an exception to this rule: as Twitter users watch a #hashtag stream, they will see tweets from anywhere in the network. However, these tags are by default filtered by Twitter to show only top tweets, so the user must actively request an unfiltered view of the #hashtag stream.

Link Sharing. Previous studies suggest that the majority of retweets contain links [4], a far larger proportion than is representative of original tweets. Two popular types of content retweeted are links to social media content sites such as Flickr, yfrog, FaceBook or YouTube and links to web-based content sourced from traditional media institutions such as online newspaper articles.

The extent to which discourse on general social media and on Twitter, in particular, is driven by individual initiative or emerging consensus versus broader dissemination through traditional media is an open question. Zhao and Jiang [5] performed an empirical study comparing content coverage on

Twitter with that of the *New York Times* (NYT) on a total of 96 million tweets, which they compared against approximately 12,000 contemporaneous news articles crawled from the NYT. The topics of each were identified in a semi-automated manner, loosely categorized into event-oriented, entity-oriented and long-standing topics. Their findings unsurprisingly confirm that event-oriented topics constitute a larger proportion of NYT than they do for Twitter. But they also found that event-oriented topics are more actively spread through retweeting, particularly when significant world news such as natural disasters was the topic. Zhao and Jiang concluded that Twitter users are more likely to spread world news than other types of information.

Social Networks and the Flash Mob. Social media, in general, and Twitter, in particular, have long been associated with spontaneously arising behavior such as the flash mob. The term *flash mob* applies to a decentralized group activity that typically has a performance aspect; for example, spontaneous group outbursts of dancing; a crowd converging on a public space to stand in motionless silence for a predefined time before moving on as though nothing had happened; or light-saber battles in a shopping center. Gore [6] traces the history of the flash mob to precursors in the 1960s and identifies as well similar concepts such as science-fiction author Larry Niven's "flash crowd." She describes it as a concept in flux, moving away from an initial definition as "pointless" performance activity or "gratuitous acts of fun" to become associated instead with "celebration, political activism and/or commercial advertisement" [6, pp.126-127]. Nonetheless, the unthreatening organized spontaneity of the flash mob is a benign form of "swarming," which is defined by White as "the unexpected gathering of large numbers of people in particular public locales" [7, p. 321]. Flash mobs usually depend on the same technical communication links that are said to facilitate more sinister applications, such as organized gate-crashing and, indeed, riots.

The London Riots. During the week of August 6, 2011, a series of riots occurred in England. The point of origin of the disturbance was Tottenham (in the greater London area), following a fatal shooting that occurred with unclear circumstances. This death sparked a peaceful protest. However, rioting began shortly thereafter and spread through London, nearby towns

and cities and eventually to several of England's larger cities, including Birmingham, Nottingham, Leicester, Wolverhampton, Liverpool, Manchester, Bristol and others. During this time period, 3,443 crimes related to the riots were recorded across London. Over a thousand people were charged as a result. The sentences handed out divided public opinion, with some seeing the penalties as unusually severe and others describing them as too light. It is clear that the impact of these events on Twitter users provides only a partial accounting of them. Nonetheless, here we will present the events from the Twitter perspective, but as with any social medium it is important to note the limitations of the source.

As examples, here are some notable events or movements:

- Vigilante groups formed. A notable example: a group of Millwall Football Club supporters, themselves typically associated in the media with hooliganism, controversially took to the streets as vigilantes.
- Large groups of volunteers offer to clean up the damage, arriving at various locations armed with cleaning materials. Various appellations were used for these efforts; following the widespread distribution of an image showing a crowd carrying brooms, many used the term *broom army*. Another popular term, and hence #hashtag, was #riotwombles (#wombles), a nostalgic term that represented the Wombles, which are characters in a cult-classic BBC television series, itself based on a set of children's novels. They are a species of furry characters who live under Wimbledon Common in London and occupy themselves by cleaning the common of rubbish, which they then recycle for their own purposes.
- The hashtag #OperationCupOfTea, which called for participants to "stay in and drink tea" (that is, establish a self-imposed curfew). An alternative interpretation related to another symbolic image: a police officer using an upturned riot shield as a tea tray.
- Campaigns were set up to help individuals particularly affected by events; several such initiatives raised around £30,000 (\$45,000) each for people including 89-year-old barber Aaron Biber and shopkeeper Siva Kandiah to pay for repairs to commercial premises and contents. Perhaps

the highest-profile case was that of the #somethingniceforashraf appeal for Ashraf Haziq, a Malaysian student who, having been injured, was then robbed by a group of youths under the guise of rescuing him.

It is perhaps relevant to note that during the London riots, two men were given four-year jail sentences for inciting to riot by creating Facebook pages or Facebook events, although in neither case did anyone turn up for them [8] a fact attributed by some to preventative action by the police. As others have found in less extreme circumstances, the use of social media as an organizational tool does not necessarily lead to a demonstrable result. Any flash mob event, even the most harmless, may sound compelling online, perhaps even receiving considerable attention on Facebook or other sources. Flash mobs may fail for any of several reasons ranging from lack of interest, to poor timing, coming to the attention of the authorities – thus discouraging potential participants – or an intimidating/inappropriate choice of venue [9, 10].

Here we look at the following questions, reviewing publicly visible evidence appearing within our dataset, to explore how Twitter was used during the riots:

- 1) Is there evidence to suggest that Twitter was used as an organizational tool during the riots, and what uses were made of it?
- 2) What can we learn from retweets? Do tweets primarily influence individuals toward new concepts, or are retweets performed primarily when the tweet confirms existing beliefs?
- 3) What uses/analysis might be made of real-time data from Twitter during events such as these?

As a consequence of the sudden ascendance of Twitter in the media environment, researchers have come to view it and similar tools as a useful resource susceptible to a broad variety of possible uses and analytical methods. Sentiment analysis, for example, may be applied for opinion mining, which, in turn, may be used by businesses in order to gather opinions about a product or the public image of an individual or establishment or for political purposes to establish the popularity of policies or proposals [11]. However, it is not an easy task. The enforced brevity of messages (140 characters or less) and the informal and specialized language make established approaches, such as

part-of-speech tagging and sentiment lexicons, less effective [12]. Message content may be subjective or objective, informational or opinionated; Barbosa & Feng [13] find that the most prolific tweeters frequently are advertising a product or service.

Methodology

Tweets were harvested and stored in a local database. Twitter wasn't always available, possibly due to server load, so coverage is incomplete. The dataset includes the August 9-11, approximately half of the disturbance, as well as its aftermath. Coverage is limited to tweets mentioning selected hashtags, notably #londonriot and #riotcleanup.

Following the harvesting process, we began by taking the following steps:

1. Identifying retweets and duplicate tweets
2. Identifying references to other participants in the social network, enabling a graph of participation to be built
3. Identifying repeated phrases, loosely describable as noun phrases plus relevant, usually adjectival, modifiers (*police car, burning police car, vigilant people, long night,...*).
4. Identifying individuals and locations (*David Cameron, Clapham, ...*)
5. Identifying URLs and preprocessing them in order to identify their target URLs (if the link goes through a URL redirection service).

The quantity of data involved in this sort of study is large enough to make automated (or semi-automated) analysis highly desirable, provided the methods used yield sufficiently high-quality data. Manual data analysis must be performed where necessary in order to check the accuracy of these methods.

We used natural language processing to identify interesting terms, indexing tweets according to the occurrence of these phrases. This enables us to build up frequency tables of interesting phrases, allowing us to identify popular topics and change in topics over time. The data includes a great deal of noise, but there are various ways of mitigating this problem – machine learning methods are particularly useful, using existing knowledge to clean up results and improve accuracy. One task that does present a

problem is identifying individuals by nickname or partial reference, although working from an ontology can help with this challenge.

There are similar difficulties of ambiguity in identifying locations. A contemporary gazetteer (a location database designed to support the extraction and resolving of place names from texts) may be used for this purpose. However, the success of this tool is at best partial, since there are many examples of contextually bound information, such as “on the street outside” or “near my office,” terms that require more contextual information to resolve successfully.

URL shortening services (that is, services that provide a short URL, or web link, to be used in place of the original link) typically provide a redirect instruction to the browser, pointing it to the intended destination URL. We found well over two dozen URL shortening services in use, including “vanity” services from many newspapers and news services, including British, American and European services. Vanity services are designed in part as a utility and in large part for the purposes of branding. For each link, we identified both initial and destination (resolved) URLs.

Results

Six hundred thousand tweets were stored and analyzed as described above for topic, content and reference.

Topic of Tweets. Our findings indicate that, twinned with a supervised learning algorithm such as a Bayes filter (a system trained on an existing reference data set) in order to classify types of entity, it is plausible to identify individuals (David Cameron, Boris Johnson), named events (England v Holland match), times/dates (tomorrow night), organizations (Sky News) and collective identifiers (London rioters, police officers). It is also possible to identify theme-relevant compound terms such as *rubber bullets* and *water cannon* (proposed for use, but not employed during the disturbances) as well as *social media* and *Facebook page*, terms that appear extensively as a result of discussion in the traditional and new media about the role of social media in organizing and propagating the riots.

Proportion and Content of Retweets. Retweeting is commonplace; slightly under half (over 48%) of the tweets identified contained a retweet. Slightly

over a quarter of non-retweets (original tweets) contained one or more links. Of retweets, 45% in our dataset contain one or more links (an actionable URL). These refer to a variety of sources, notably traditional media web presences and social media sites, including social networking sites such as Facebook, image hosting sites such as yfrog and Flickr, and others.

Twitter and Media. A complete review of the most popular resources falls outside the scope of this paper, although we intend to review them in detail in future. The evidence shows us clearly that many of the individuals tweeting on the riots draw inspiration from news articles or directly make reference to news articles. News sources referenced include British newspapers, TV stations and websites; international sources are also discussed, including central European sources and North American sources among others.

The most common phrases visible in the dataset appear to be driven by traditional media. While the danger zones for riots in London, Birmingham, Manchester and Bristol are represented in the data, they appear with a relatively low frequency compared to themes arising from media coverage on topics related to the riots.

The influence of the classical media coverage is clearly visible from the extracted noun phrases: for example, *water cannons*, *rubber bullets*, the *England vs Holland Match tomorrow night* and so forth. From the evidence collected, it appears that most people were either tweeting or retweeting news related to the riots. By contrast, original information or opinion appeared in a minority of tweets.


Social media as an information source. Analysis of the most popular links suggests that a large proportion of the most frequently mentioned links are themselves products of social media – images published online through yfrog, for example. Interestingly, the single most popular sites referenced are images rather than textual resources; for example, there is a famous image from this time, showing a sign from the front door of a Subway restaurant containing the handwritten text, “Due to the imminent collapse of society, we regret to announce we are closing at 6pm tonight.” Each of these very popular images is separately identified by a large number of web links – several different URL shortcuts have been used by different twitter users.

There are also many resources from Facebook, blog sites such as Tumblr and others which have become major parts of various campaigns such as the money-raising campaigns mentioned earlier. Another popular set of resources is the e-petitions website belonging to the UK government, which was quickly populated with a set of demands, requests and suggestions relating to the riots, such as “Convicted London rioters should lose all benefits.”

Identifying New Information. Through detailed content analysis (semantic analysis of part-of-speech tagged data such as searching for specific types of event, categories of incident and so forth) it is possible to extract various examples of reports of specific events – fires, crimes, police presence and so forth. However, Twitter seems to be less effective for crowd-sourcing than might be expected, due to the huge RT (retweet) echo effect from those who are not directly involved, most of whom appear to be reacting to media coverage rather than directly reporting individual experience.

Performance Tweets. Of apparently suspicious tweets, there are a number that riff on current events for comedic value. For example, a tweet that ostensibly advertised a Craigslist London advertisement for 40 16GB iPhone 4 units, “brand new and sealed,” in fact sarcastically parodied the original advertiser’s response to queries regarding the devices’ authenticity: “Honest, I bought 40 iPhone 4s for friends and family and they just don’t want them.” Some can be identified through telltales such as references to memes

FIGURE 2. The overall popularity of hashtags through the corpus, plotted on a log-log scale



(loosely definable as “widely propagated catchy ideas” [14]).

As is the case in sentiment analysis in general, sarcasm and metaphor are stumbling blocks. Human judgment, if supplemented with relevant contextual information and background knowledge about the community, is still not always able to tell for certain whether a tweet is satirical or intentionally inflammatory in intent; context is the key.

Hashtag Selection. The emergence of a dominant #hashtag (as seen in Figure 2) for a channel is not a simple matter of nomination and acceptance. Sometimes it is a collaborative process, sometimes it is competitive and often the process seems to involve elements of both.


It could be described as a special case of what Cattuto et al. [15] call “semiotic dynamics,” meaning the process of developing and sharing a name for something, with the practical aim of using it to talk

about the object in question. Leaving aside the details, the development of a term from a temporary label to a general convention is a feature of many parts of the social web, such as social tagging [15] and #hashtag development [16].


In Figure 3, the tags are associated with actions that individuals can take as part of their beliefs or agree to promote for the original tweeter. The chart shows approximate numbers, estimated through Google’s site-search function. An interesting result seen in Figure 3 is the way in which apparently duplicate terms can coexist; another is the use of idiomatic phrases. Many of these phrases are in-jokes of one kind or another and are not very accessible to those who are not familiar with British humor.

Our methodology was based on the assumption that popular resources would be referenced under various names, and Figure 4 shows that such variety is indeed the case. On the whole, it seems that popular resources are referenced under many names. This practice has several effects, one being that it is difficult without individually checking out the destination of every

FIGURE 3. Popularity of tags associated with actions related to the London riots



TONKIN, PFEIFFER and TOURTE, continued

FIGURE 4. Multiplication of referenced resources given unique identifiers


link to find out precisely how many people are talking about the same thing – indeed it is quite possible for many people to talk about the same web page without knowing it! Another is the extreme dependency on URL shorteners. Their use may make it difficult for people to understand the context of archived conversations, since, if these shortened links no longer work, people will no longer be able to see the resources being discussed. Discontinuance is a problem not only because of the possibility that the resource itself has been taken offline, but also because the URL shortener may have been discontinued as a service.

Hashtag Adoption and Individual Twitter and Media Profile. It is often difficult to identify the source of ideas, including #hashtags. Retweets of an individual tweet do not necessarily identify the originator of the idea; instead, they usually identify the individual who formulated the tweet in its eventual, popular form. To explore the question of whether there is any correlation between the popularity of the hashtag and the new or traditional media profile of its supporters or originators, we extracted the top 10 entirely unaltered

retweets (retweets that contain only the body of the original tweet, along with the retweet markup). These examples were chosen because there is no obvious mechanism to analyze the sentiment behind the decision to retweet; that is, there is no contextual information available on which to base a judgment. For the sake of argument, we assume that in the majority of cases the individual agrees with the content of the message.

Statistics from these retweets show that of the top 10 retweets in this category, the majority – and four of the top five – originate from celebrities and media figures. The most popular originated with Piers Morgan, a former newspaper editor and now a television presenter who has recently moved to CNN to replace Larry King with an evening interview program. The originator and driving force behind #OperationCupOfTea was Sam Pepper, a contestant on the reality television show, Big Brother. Others included a crime journalist; a musician; actor, writer and comedian Simon Pegg; and former pop musician turned television presenter and physics professor, Brian Cox.


Figure 5 reviews the top 25 unaltered retweets. It compares the number

FIGURE 5. Retweets and celebrity individuals

of followers attracted to a given account with the number of retweets received by a popular tweet originating from that account. The result suggests that there is a visible variation in the ratio between the number of followers and the number of retweets received; graphing this information shows a fairly coherent group of non-celebrity Twitter users, which appear in a different region of the graph from the broader range of celebrity users. The usefulness of this observation depends on how well it generalizes to a broader dataset, which we will investigate as part of our future work.

Conceptual Ontology. It is interesting not only to look at the content, what is being said within the tweets, but to analyze the structure of the tweet #hashtags. Viewing the #hashtags over time, these tags can be placed into a

FIGURE 6. A conceptual ontology of popular #hashtags associated with the London riots



hierarchy of parent/child relationships depicting how the information of the tweets' #hashtags are being passed along. This hierarchy is referred to as a *conceptual ontology* [17], where the

concepts in the hierarchy are actually the #hashtags themselves.

In Figure 6, the #hashtags from the top 10 retweets were analyzed over time to see which #hashtags appeared first, next and so forth, such that, later tweets, with other #hashtags (such as siblings), can also be seen. In these top 10 tweets there were a total of 6 #hashtags used. As we have already seen

earlier in this article, #londonriots was the first and number one #hashtag of the retweets and appeared the most often; therefore, this #hashtag will be put at the “root” (top) of the hierarchy indicating that all other #hashtags have something to do with this one. The next three #hashtags appear in tweets as siblings; that is, they are not emanating from each other, but are equals. Therefore, in the ontology they will be at the same level. In tweets about the Blackberry Messaging Service (BBM or BMM) it is interesting to see that tweets using the #hashtag referencing the #BBM (support of) and tweets referencing #BlockBBM (block out BBM) are seen as siblings in the hierarchy and not as parent/child. We also see that out of the #BMM hashtag grew the #OperationCupOfTea, but not out of #BlockBMM. This indicates that people in support of the BBM were also in support of #OperationCupOfTea (remember that is people with their own imposed curfew), while people not in support of the BBM (#BlockBMM) were not retweeting the #OperationCupOfTea hashtag. The hierarchy, thirdly, indicates that #wombles hashtags grew out of the #riotcleanup people, but not necessarily the people that were following #BBM or #BlockBMM hashtags. More analysis of the conceptual structure of the #hashtags' ontology will appear in future work.

Conclusion

Looking at data from Twitter may be challenging, but it is also worthwhile – it frequently reveals fascinating trends. From the data that we were able to retrieve from Twitter on the London riots, three aspects of tweeting became clear: 1) people support and retweet conversations that support their beliefs; 2) tweets that are off topic and have unrelated information die on their own; and 3) tweets that are made by popular or newsworthy people are more supported (retweeted), in general, than non-notable people. There's little data to support the idea that the service was widely used for inappropriate uses – tweets of this nature generally either die out without being retweeted or are retweeted to “name and shame” the offender. ■

Resources on next page

Resources Mentioned in the Article

- [1] Suh, B., Hong, L., Pirolli, P., & Chi, E. H. (2010). Want to be retweeted? Large scale analytics on factors impacting retweet in Twitter network. In *Second IEEE International Conference on Social Computing (SocialCom'10)* (pp.177-18).
- [2] Boyd, D., Golder, S., & Lotan, G. (2010). Tweet, tweet, retweet: Conversational aspects of retweeting on Twitter. *Proceedings of the 43rd Hawaii Conference on System Sciences [HICSS], 2010* [pp. 1-10]. Newton, MA: IEEE Computer Society.
- [3] Liu, H. (2007). Social network profiles as taste performances. *Journal of Computer-Mediated Communication*, 13(1), article 13. Retrieved November 10, 2011, from <http://jcmc.indiana.edu/vol13/issue1/liu.html>.
- [4] Zarrella, D. (2009). The linguistics of retweets [Blog post]. Retrieved November 11, 2011, from <http://danzarella.com/retweet-linguistics.html>.
- [5] Zhao, W. X., & Jiang, J. (2011). *An empirical comparison of topics in twitter and traditional media*. Singapore Management University School of Information Systems Technical paper series. Retrieved November 10, 2011, from [www.mysmu.edu/faculty/jingjiang/papers/TechReport\(Zhao2011\).pdf](http://www.mysmu.edu/faculty/jingjiang/papers/TechReport(Zhao2011).pdf).
- [6] Gore, G. (2010). Flash mob dance and the territorialisation of urban movement. *Anthropological Notebooks* 16(3), 125-131. Retrieved November 10, 2011, from www.drustvo-antropologov.si/AN/PDF/2010_3/Anthropological_Notebooks_XVI_3_Gore.pdf.
- [7] White, R. (October 2006). Swarming and the dynamics of group violence. *Trends & Issues in Crime and Criminal Justice*, 326. Retrieved November 10, 2011, from www.aic.gov.au/en/publications/current%20series/tandi/321-340/tandi326.aspx.
- [8] Gross, M. (September 2011). Why do people riot? *Current Biology*, 21(18), R673-R676. Retrieved November 11, 2011, from <http://download.cell.com/current-biology/pdf/PIIS0960982211010177.pdf?intermediate=true>.
- [9] Walker, R. (2011). *Eight is not enough: A historical, cultural, and philosophical analysis of the flash mob*. Unpublished doctoral dissertation, Louisiana State University. Retrieved November 10, 2011, from <http://etd.lsu.edu/docs/available/etd-04112011-102646/>.
- [10] Wasik, B. (2006). My crowd: Or, Phase 5: A report from the inventor of the flash mob. *Harper's Magazine*, 312(1870), 56-66.
- [11] Pak, A., & Paroubek, P. (2010). Twitter as a corpus for sentiment analysis and opinion mining. In *Proceedings of the 7th International Conference on Language Resources and Evaluation (LREC2010)*. Retrieved November 10, 2011, from www.lrec-conf.org/proceedings/lrec2010/index.html.
- [12] Kouloumpis, E. Wilson, T., & Moore, J. (2011). Twitter sentiment analysis: The good the bad and the OMG! In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*. Retrieved November 10, 2011, from www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2857/3251.
- [13] Barbosa, L. & Feng, J. (2010). Robust sentiment detection on Twitter from biased and noisy data. In *Proceedings of the 23rd Annual Conference on Computational Linguistics (Coling 2010)*. Retrieved November 10, 2011, from www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2857/3251.
- [14] Knobel, M., & Lankshear, C. (2007). Online memes, affinities, and cultural production. In *A new literacies sampler* (pp. 199-227). New York: Peter Lang.
- [15] Cattuto C, L.V.& Pietronero L. (2007). Semiotic dynamics and collaborative tagging. In *Proceedings of the National Academy of Sciences*, 104(5), 1461-1464.
- [16] Lindgren, S., & Lundstrom, R. (September 2011). Pirate culture and hacktivist mobilization: The cultural and social protocols of #WikiLeaks on Twitter. *New Media and Society*, 13(6), 999-1018.
- [17] Sowa, J. (1984). *Conceptual structures: Information processing in mind and machine*. Reading, MA: Addison-Wesley.