A Literature Review of Recent Microblogging Developments

Technical Report

Marc Cheong, Sid Ray

Clayton School of Information Technology,
Monash University
Victoria, 3800
Australia

{marc.cheong,sid.ray}@monash.edu

Abstract. In this paper, we review Twitter and microblog research from 2009–2010 including but not limited to fields such as visualization, HCI, machine learning, computer-supported collaborative work, online social networks, media and social studies. We also propose a categorization scheme for the different aspects of research in microblogging, while improving on the categorization scheme found in earlier surveys.

Table of Contents

1 Introduction

Twitter is a microblogging service that has fast become commonplace in both our daily lives, and also as a subject in itself in academic research. From its beginnings as a novelty 'status updating' service with the tagline *"What are you doing?"*, it has evolved into its current form which is more focused on mirroring the goings-on in the real world, with the new tagline *"Whats happening?"*. Twitter allows users and organizations to publish messages, communicate with other users in realtime, and even use scripted programs or bots to perform tasks which utilize delivery of short messages.

Twitter and related microblogging technologies have been given much attention in multiple disciplines since its inception in 2006. In Cheong and Lee's [1] earlier literature review, we find that the pioneering academic papers focusing on Twitter became available circa 2008-2009, and literature has continued to mushroom even as the aforementioned paper was still being written.

With the increasing number of research work dealing with Twitter, the fields of study and application of Twitter research has also broadened; Cheong and Lee originally reviewed literature from disciplines such as "social media and the Web, ... anthropology, computer-human interaction, data mining, knowledge discovery, visualization" [1] to name a few. This list has significantly expanded; as from our findings, new areas of study hitherto not associated with microblogs (such as terrorism informatics, user modeling and personalization, online security, spam detection, and information streaming) have been studied with Twitter and microblogs as their primary focus.

The literature surveyed and critiqued in this review comes primarily from conferences or publications dealing either directly with microblogging/Twitter; on applications of microblog technologies in research areas such as communications, information systems, social networks, and computer-human interaction (to name a few); or for research areas which are of potential to be applied to microblogging in future research: visualization for instance. This paper also reviews all the submissions to the *CHI 2010 Workshop on Microblogging*, which to the authors' belief, is the first workshop dedicated solely to emerging research in microblogs with a particular focus on Twitter.

2 Scope and categorization

We adopt the classification methodology used in the earlier review [1] which identified two broad concepts or domains (that are independent of one another) found in Twitter and microblogging services. These are:

1. **the user domain**: properties exhibited by a user in a microblogging environment accessible via the Twitter API; this includes statistics such as tweet count, account age, user customization and so forth, which allows us to study the human factors behind microblogs.
2. **the message domain**: properties exhibited by a single message composed by a Twitter user. The raw data extractable from Twitter API for example

includes the message content, software client used, timestamp, geo-location properties, and any embedded content; which embodies all the characteristics of an individual message.

The classification used in [1] has since been backed by independent research, e.g. by Cormode *et al.* [2] who studied the modeling and measurement of online social media – Twitter inclusive. Their research indicated that "there are actually two central objects [in Twitter]: the user and the tweet itself" [2]. Hence, in this paper, our findings from the state-of-the-art research will relate to either one or both of the user and message domains.

The categorization scheme for research reviewed in [1] is used to categorize research found in this paper, albeit with certain adaptations to accommodate for new types of research that we have encountered; this includes minor changes to subsections to better reflect the evolving categories of microblog research since [1] was authored.

3 Exploratory Studies

Besides the already-popular body of work on the exploration of the Twitter 'ecosystem' [3–5], several new papers have been written that could contribute to knowledge of this subject, from the combined perspective of both Twitter users and messages.

3.1 Properties, Emergent Features, and Evolution

Humphreys [6] has drawn parallels of microblogging practices with handwritten diaries from the eighteenth/nineteenth centuries. In his findings, four parallels exist between current microblogging practices, such as those found on Twitter, to the diaries of yore: (1) they are both semi-public in nature; (2) they both introspectively chronicle activities and mundane day-to-day trivium; (3) they are both in the narrative form (e.g. entries which discuss upon tragedies); and lastly (4) diary entries are rather short due to limitations of space, similar to microblogs' 140-character constraints. One of the obvious differences is that the diaries lack social interaction; in the sense that technological advancements allow microblog users to mutually 'follow' one another.

On the subject of content, Petrovic *et al.* [7] have made available the *Edinburgh Twitter Corpus* of approximately 96.3 million Twitter messages for study and analysis; however as of time of writing it has since been removed due to legal issues with Twitter Inc. However, their paper [7] summarized the most popular users and most common content on Twitter throughout their survey (November 2009–February 2010). In their corpus they have identified that:

- six out of ten of the top followed users were musicians or singers (e.g. `@justinbieber`) building a fan base on Twitter

- the top hashtagged topics on Twitter dealt with musical-based memes (e.g. `#nowplaying`) where users discuss music they are currently playing; political memes (e.g. `#tcot` or top conservative politicians); tags indicating Facebook co-usage with Twitter; and 'just-for-fun' Internet memes.
- 80% of the top Twitter client applications are the web interface, UberTwitter (indicating a high mobile device usage), and TweetDeck (indicating usage of 3rd party applications to improve Twitter user experience

Kwak *et al.* [8] has recently performed a comprehensive study of Twitter users and messages to answer the fundamental question of whether Twitter is a "social network, or a news media". The authors claimed to have surveyed almost the entirety of the Twitter user space, with approximately 41.7 million messages, 1.47 billion friend/follow links, 106 million tweets. This is achieved via a cluster of 20 machines, each limited to 10k API requests per hour to avoid violating Twitter terms of service (their harvesting includes 10 Trending Topics every 5 minutes; and 1500 tweets are harvested in the same period of time). A summary of their findings are as follows:

1. **Spam identification:** They have devised a sanitization method to remove likely spammers and 'noise' from their data: users with account ages of less than one day, and tweets which contain 3 or more Trending Topic mentions are likely to be spam capitalizing on the discussion of a particular topic. In their dataset, 20.2 million such messages and 1.9 million such users have been flagged as such. Also, irregularities in the distribution of number of followers versus number of authored tweets for a given user can be conjectured as being likely to be a spammer account.
2. **User network topology:** Based on topological analysis, they have identified that top users with more than 10k followers are only celebrities or politicians; and those with more than 1 million followers are mainly celebrities and media outlets (such as CNN). The majority of the users with less than 10 followers never contributed a single tweet, while there are also users who "tweet far more than expected from the number of followers" [8] (judging by their averages). Another finding is that approximately 67.6% users are not *reciprocal friends*[1] with either of their followed users – Kwak *et al.* [8] suggest that such users are 'consumers' of Twitter as an information source. Finally, Kwak *et al.* studied the topology in terms of user homophily and determined that users with 1000 followers or less are geographically close (based on the time zone property of their user accounts) to their reciprocal friends and have a similar measure of popularity (in terms of number of followers, not necessarily reciprocal). Their findings from this part of the survey concluded that Twitter diverges from normal online social network (OSN) traits where: it is not power-law distributed, users have a short degree of separation on average, and not all friend/follow links between users are reciprocal.

[1] *reciprocal friends*: mutually followed by the user they are following; this term will be use throughout this paper to define such a bidirectional user relationship on Twitter.

3. **User rankings:** The authors first applied the PageRank algorithm to their user graph to study user ranking from a network perspective ("propagation of influence"). Next they performed an analysis of user popularity in terms of number of retweets (RTs); top retweeted users are politicians and musicians, complementing findings from [7]. Interestingly, when both sets of users are compared using a generalized Kendall's tau, they found that there is a discrepancy between the number of followers and the popularity of RTs, bringing "a new perspective in influence" in terms of Twitter [8].
4. **Trending topics:** This study surveyed 4,266 unique Twitter trending topics (TTs), covering major issues from June to September 2009, complementing other studies on trending topics on Twitter during the time period such as [9]). Several observations have been made with regards to this dataset of topics. Firstly, TTs have been compared to Google Trends which reflect trending search terms: the similarities in terms of textual substring matches is insignificant (a mere 3.6%). However, more importantly, TTs are found to be more persistent in comparison with Google Trends: TTs have only 72% new topics, 23% less than Google. Another finding which is of benefit to further research on trends is Kwak *et al.*'s classification of TTs in the context of tweets to be either: a singleton (normal tweet without mention of users or retweets); mentions (with notation `@user`, but not link to a previous message); replies (mentions which are linked to a previous message); and retweets (with the retweet keyword `RT`). Generally they have found that "singletons are most common, followed by replies and retweets" [8]. A more concise discussion on trending behavior of topics will be discussed in Section 4.5.
5. **Retweeting behavior:** The authors found out that the "distribution of the users in a retweet tree [graph representation of retweets] follows [the] power-law distribution," despite the user friend-follow connections not adhering to it. The median of a message retweet is less than an hour, and that retweets are generally diffused rapidly after the second level of retweets (i.e. a retweeted message itself being retweeted). Fifty percent of retweets happen within the first hour; 75% within a day; and 90% within a month. Favoritism in retweets is evident in that a limited subset of a users' followers actually retweet the message; and that despite the follower count of the original tweet's author, the tweet is "likely to reach a certain number of audience, once the user's tweet starts [being retweeted]."

To conclude this section on properties and emergent features, we cover Yoshida *et al.*'s [10] paper on the frequency of URLs in tweets, and the proliferation of bot-generated content. By using data from both the Twitter public timeline and a custom dataset incorporating Japanese links, they obtain approximately 19.7 million tweets containing 20 million URL mentions; after filtering for nonexistent webpages and duplicates, they obtain 12.7 million tweets in total. By analyzing the source (software client) used to contribute tweets, they observe that the web interface, Twitter API-based bespoke programs, and RSS feeds top the list. The authors also analyzed the tweets for the proportion of software bot-generated tweets, and found a 35.02% ratio in such automatic postings in their Japanese link dataset (9.01% for their public timeline samples for comparison). One key distinguishing factor between bot-posted tweets and human-posted ones is that after excluding the URL string in a tweet, bot postings have a higher average string length (48.89 characters) compared to human postings (41.51 characters). For bots which usually truncate the length of posts, the distribution of tweet lengths have a high peak at the right side of the graph (about 110-120 characters out of the theoretical maximum of 140 *sans* the URL string). By contrast, however, retweets for human posts are higher (12.55%) compared to bot posts (a mere 1.44%). Content-wise, the most frequently shared URLs typically consist of photo-sharing websites (the highest URL count, found in 11.9% of human-posted tweets), media sharing, and news agencies (which form the bulk of retweeted content).

3.2 Measurement and Statistics

There have been several papers reporting on current measurements and statistics of Twitter in its entirety, from both the user and message domains, which we shall discuss in the current section.

Pear Analytics [11] has published statistical reports on Twitter activity, growth and similar metrics. Vital points in their August 2009 report (the most current, as of time of writing) include:

1. Users per month (USA): 27 million (June 2009)
2. **Demographics**:
   - 55% female
   - 43% in the (18-34) age bracket
   - 78% users of Caucasian descent
   - average annual income of US$30-60k.
3. **Uneven distribution**: 1% of top users contribute 35% of visits; with 72% ‘passers-by’ as opposed to 27% regular users
4. **Types of posts**: Mainstream news, spam, self-promotion of businesses, babble (everyday trivium; this being 40.55% of total posts), conversations (37.55% of posts), and pass-along messages (retweets)
5. **Frequency by time of day**: morning for retweets, end of work day for conversations, midday for news and babble (with spam and promotions peaking twice in the workday).
6. **Frequency by day**: early week for retweets, mid-week for news and conversation, end of work week for spam and babble

The abovementioned statistics (especially the frequency distributions) – in our opinion – approximate those of a typical office-worker; this reflects of the demographics of users tending towards those in their twenties which contrast against users of other online social networks who are younger (e.g. as found in [12, 13]). Another concern found in this report that partially backs this claim is that Twitter adoption is less frequent among younger people as it is deemed “not safe” due to the lack of "ability to select who they want to connect to" [11] as opposed to de facto social networks such as Facebook and Myspace which allows the use to control the degree of information shared (further discussion in Section 6.4).

Krishnamurthy [14], in his discussion of challenges in measuring online social networks [2], has characterized Twitter in the Web 2.0 OSN context as a *micro-OSN* by defining several properties of interest, which we detail and expand upon (in parentheses) in Table 1.

Feature class	**Feature**	**Twitter's implementation**
Profile details	Age/Gender	None (can be predicted/conjectured in research)
	Location	User text or GPS coordinates
	'Testimonials'	None (closest is a summary in user timeline)
Connectivity	Friends	Non-mutual 'followers' relationship
	Subscriptions	Non-mutual 'friends' (following) relationship
	Groups	Lists API
Content	Main	Microblog entries: 140-character tweets
	Other	Linked URLs in text
	Tagging	Hashtags
	Friends only	None (either completely private or completely public)
	Comments	Tweet reply, using `@user` notation
	Editing	None (can only author/delete)
	Rating	None (closest is implicitly 'favoriting' or retweeting)

Table 1. Twitter features as a Web 2.0 micro-OSN

Krishnamurthy [14] also mentioned several unique properties of Twitter that aren't obvious in other studies previously covered: presence of cultural bias (as a few countries dominate in terms of geographical Twitter user distribution), language bias exists ("Japanese users tweet in Kanji and do not have many English speakers as followers" [14]), and the presence of Intra-European cliques in terms of user communication patterns.

On a related topic, Cormode *et al.* [2] presented a 'manifesto' for modelling and measurement of social media, and discussed several challenges and observations for Twitter as a micro-OSN. Twitter has unique relationships between its 'entities' (e.g. friends, followers, hashtags, replies, retweets) which is different from the classical model of networks consisting of nodes and edges. Also, they discuss challenges on Twitter research in general, such as the various ways to access data (e.g. API versus HTML scraping versus traffic sniffing) and sampling methodology (e.g. random node identification versus crawling within a boundary). Several issues regarding quality of data have also been identified in Twitter: presence of dormant users distorting Twitter friend/follow network properties; constant redesign (such as the newly-introduced Lists feature as of end 2009 and

[2] *OSN*: online social networks; Twitter is classified as a micro-OSN because it has a limited subset of features of an OSN

changes in API); and usage of existing features "in ways that are unanticipated". The most striking example of the latter point comes from the "formalization of previously unsupported conventions adopted organically by Twitter users" [2] such as retweets and hashtags; another example is the usage of celebrity pages in place of profile URLs. Their case study involving Twitter has formalized the notion of using the concepts of 'users' and 'messages' as "central objects" in Twitter [2], justifying an earlier dichotomy in [1]. As a concluding note, the authors hinted on the limitations of the Twitter API on future measurement and research of Twitter, such as the recently available Twitter stream API only making a sample of the tweets available, the difficulty of measuring the fraction of private tweets, the absence of a dedicated 'grouping' function, and the bias in observations caused by new users and spammers [2].

3.3 Friending/following, and Hashtagging Behavior

In terms of Twitter feature usage, behaviors of message addressivity and replying (`@user` messages) [15] and retweeting (forwarding using the `RT` convention) [16] have been discussed in the earlier paper by Cheong & Lee [1]. However, studies on other features of Twitter such as friending/following, and hashtagging were not covered. The aim of this section is to cover these features – one of which is a formalized but "previously unsupported convention" [2, 17] – to complement the earlier survey [1].

Baumer and Leis [18] studied the "genres of participation in [Twitter] *following*": they note a shift in media consumption patterns from classical blogs to microblogging; and the existence of an inherent assymmetry of interaction, from the author to the readership of his material. They classified following patterns into the two categories of *minimalists* and *zealots*:

1. **Minimalists** tend to have 10-30 close friends, use Twitter to socialize, and are normally recommended by others to join Twitter. Such participants "see Twitter as a more intimate... way of connecting" [18] with people they are interested in.
2. **Zealots** on the other hand tend to follow hundreds of people, comprising of friends/colleagues and information sources. To them, Twitter is mainly for "professional and information-seeking purposes" [18], and generally adopt Twitter by themselves without being recommended by others. Also, they tend to experiment with Twitter from a variety of software clients.

The authors in their qualitative user study revealed that despite the differences in user behavior, both zealots and minimalists never use the Twitter trending topics feature, and that there seems to be a 'lag time' after signing up to Twitter before their usage became regular [18].

A similar study by Heil and Piskorski [19] revealed Twitter users are more likely to follow others from the same gender. Male users are twice likelier to follow other male users; while females are 1.25 times more likely to follow other female users. They also found that 10% of top Twitter users contribute to about

90% of the content, which starkly contrasts with a typical online social network (which should have the top 10% users contributing 30% content instead). They surmise that Twitter's friend/follow distribution is more of a 'one-way publishing' pattern rather than a social network connecting peers, which is similar to the findings in [8]. Studies in political affiliation in the US by Metaxas & Mustafaraj [20] also revealed that users tend to follow similar users (by political orientation in their study), as the political affiliations of the top 200 users in their dataset is correctly predicted about 98% of the time, simply by observing their *following* habits.

Huang *et al.* [21] performed a study on conversational hashtagging on Twitter. While their study deals with interpretation and statistical trend analysis of hashtags, the quantitative trend analysis aspect will be covered in Section 4.5. With respect to trending topics on Twitter (which, interestingly, was not available at the beginning of Twitter's launch) Huang *et al.* remarked that "the act of tagging a tweet increased the likelihood of a tweet being [collated and] displayed in a group of tweets on a trending topic" [21]. Hashtagging is an emergent behavior noticed among users as a means to tag messages – without intervention from Twitter staff – which is a form of "conversational tagging", where the tag itself "is an important part of the message" [21] as opposed to merely describing a message. Hashtags also sometimes turn into emergent *micro-memes*, in that users are more inclined to comment or share their views/commentary about a hashtag (topic) only *after* seeing that a particular hashtag has trended.

4 Information Spread and Self-organization

4.1 Twitter in Crisis and Convergence Events

In the original literature review [1], crisis (e.g. natural disasters) and convergence events (e.g. US political conventions) have been studied in the combined perspective of Twitter users and messages: message spread and flow, and characteristics of the Twitter user network [22, 23]. Since then, there has been more research that dealt with similar events from an international perspective, which would improve existing knowledge of the subject.

Convergence Events: Activism and Democracy This subsection details the qualitative and quantitative observations of Twitter microblogging in activism and democracy-building campaigns.

Two position papers presented at the *CHI 2010 Workshop on Microblogging* by Ems [24] and Lin *et al.* [25] provide a quick overview of the effectiveness of microblogging – particularly Twitter – in activism. Ems [24] posits that Twitter is an effective tool to communicate and disseminate information by people in authoritarian regimes; and acts as a "sieve for news media outlets" by linking to other forms of media and "amplifying the distribution of other facts" for the knowledge of others [24]. By observing recent developments in the Iran election (dissemination of pictures, video, and stories for public awareness), Moldova (to facilitate organization of protests), and the G20 Summit (in the role of 'informer' to help protesters avoid the police), the paper concludes that Twitter can be a threat to authoritarian regimes, providing power to the people by helping to shape public opinion. Lin's position paper [25] focused more on activism for disability awareness, and found out that Twitter is efficient in promoting awareness by allowing the exchange of information; allowing people to 'be heard'; building an informal social network among supporters; and allow for a viral spread of information (or 'marketing' a particular campaign). Recently, we observed that Twitter is used to provide updates on the current status of whistleblowing site WikiLeaks (`@WikiLeaks`) after the recent exposé of confidential diplomatic cables.

Several of these implications of microblogging in activism will be evident in the following case studies.

1. **2009 Iranian election controversy**: Burns & Eltham [26] performed a sociological evaluation of Twitter in the perspective of both the citizen protesters and the government. The early adopters of Twitter in this situation, by leveraging Twitter to spread awareness of the situation, was able to take the issue mainstream and reaching critical mass. Real-time broadcasts of updates and an online 'green campaign' (where supporters change the color of their Twitter profile picture to green in solidarity) was able to generate awareness of the situation in Iran to a wider audience despite threats of censorship. The spreading of the video of the shooting of a young protester, Neda, allowed the "protests [to gain] a broader, sympathetic audience" [26]. However, Twitter inadvertantly "became a vector for state repression", where it was used by the Iranian Revolutionary Guard and paramilitary to "hunt down and target Iranian pro-democracy activists" [26].
2. **2009 German elections**: Jungherr [27] has observed the use of Twitter in the German *Superwahljahr 2009* and noted a 'rapid adoption' of microblogging use by politicians, parties, campaigns and supporters. It is mainly used as a tool for community building via Twitter account 'hubs' that form a focal point of discussion, e.g. the Twitter account `@teamdeutschland`, and a distribution channel for "social objects" [27]. Similar to prior findings on US political conventions [22], Twitter is identified as a backchannel for communication in the German election context.
3. **Moldovan 'Twitter Revolution'**: Serbanuta *et al.*'s preliminary study [28], although still a work in progress, has discovered several unique characteristics of Twitter activity pertaining to this event. An analysis on approximately 28.5k tweets from 1.9k users (marked with the hashtag `#pman` describing the event) provides us with quantitative statistics which allows further research into Twitter activity rate during such an political event.
   - 14.8 messages per unique user contributing to the chatter
   - 10.7% of tweets contained links to other content
   - tweets in Romanian and English were retweeted 2.6 times on average

Crisis Events: Reactions To Disaster This subsection expands and complements upon the original papers on disaster events, e.g. by Starbird *et al.* [23] and Hughes & Palen [22].

Vieweg & Starbird [29], expanding on their earlier paper [23], have presented a position paper on analysis methods and challenges for microblogging in mass emergency events. They describe several types of future research that can be conducted on Twitter data in mass emergency: uncovering geo-location and georeferencing information, studying retweeting of data as information 'churn' and at the same time an 'informal recommender' of timely crisis information, and also understand the influence of Twitter user network connections on their "message content, user stream behavior, and information spread during crises." [29].

Kireyev *et al.* [30] experimented on using topic models on the message domain to classify microblog chatter during disasters, with unique challenges such as esoteric microblogging language patterns, short message length, and locale-specific text. As their work is still experimental, no conclusive results have been reported.

Longueville *et al.* [31] applied the approach of mining tweets for spatiotemporal data to track forest fires in Marseille, France (not dissimilar to an earlier study on river flooding [23]). They obtain a data set of 313 tweets from 127 unique users during the Marseille Fire on 22 July 2009. One of the challenges faced by the authors, which make the findings in this paper rather relevant to study disaster in a worldwide context, is that there is a low proportion of tweets from French users as opposed to e.g. US users. Temporal data is harvested from message timestamps, while spatial data is obtained from implicit place names or GPS coordinates found in a tweet (or failing which, the user profile of a tweet's author). Four main research questions have been posed, in which qualitative data analyses are applied to answer each of them [31]:

1. *Twitter is an extremely fast information dissemination platform to report exceptional events*: the timeline of tweets rather accurately matches the real-world spread of the fire, with the exception of 'lag time' at the beginning of the fire.
2. *Twitter will provide accurate and useful spatiotemporal information* [as a location based social network]: location indicators such as place name mentions, hashtags of places, quantitative measurement of location/area, and user-positioning via GPS coordinates validates this hypothesis.
3. *Users use Twitter to communicate with each other in widely open conversation; as a result, it is a primary source of information from citizens*: due to the combined presence of primary information (citizen journalism), and secondary information (aggregated data from e.g. RSS feeds which are delayed offering no added value to the information stream) within tweets, future work is required to confirm his hypothesis. However, there is evidence that citizens exhibit self-organizing behavior when tweeting during emergency situations: they quickly come to the acceptance of using unique hashtags to group their conversations in context, and center their conversations on these mutually-agreed-upon hashtags.

4. *Twitter is used as information broadcasting and brokerage platform during crisis events*: the fact that up to 75% of tweets contain URLs with links to news media, and "dozens of pictures of the fires [were] taken and published" [31] supports this hypothesis.

Another study exploring the usage of Twitter users in an emergency situation – this time on the Chilean earthquake of late February 2010 – was performed by Mendoza *et al.* [32]. The study focused on tweets in the context of the Chilean earthquake, identified using the hashtag `#terremotochile` (in English: 'Chilean earthquake'); with the threefold objectives of observing dynamics of news propagation and friend/follow patterns, influence of top users in the discussion, and to distinguish rumors versus actual news spread of a disaster on Twitter. Their dataset on Santiago, Chile-based users (based on timezone) contained 716k users with approximately 4.7 million tweets. Vital statistics of their observation on the tweet/user dataset are as follows:

– The proportion of reply-based tweets is surprisingly high, at 98%.
– Almost half the users surveyed have more followers than friends they follow.
– 'Authority users' commonly have hundreds of thousands of followers; an example of this would be CNN's Twitter account.
– The difference in distribution of user activity is evident: about 52.64% of the surveyed users only wrote one tweet during the crisis event, while 11.47% of the users have more than 10 tweets during the earthquake event; the rest are non-uniformly distributed in between.
– Users with greater than 2000 followers and/or friends will have an increase by one order of magnitude of their total tweets.
– The top 20 users in terms of activity during the event are either accounts for news media, celebrities, or non-profit organizations.
– In terms of trending topics, the fraction of users contributing to them is insignificant, compared to total users tweeting during the event.

With regards to their research foci, the authors [32] observed that Twitter discussion habits on the earthquake mirror the real world significance of the event in the real world, i.e. tweets of the disaster outnumber those discussing a popular Chilean music festival at that time. In fact, when keywords are grouped in term clouds, the temporal distribution of event words by day has a high correlation with the changing real-life events as a result of the earthquake; e.g. "tsunami" which reports on events on the first day, followed by "missing people" on the next day as a consequence. Retweeting behavior exhibited by users are tree-like (c.f. Kwak *et al.* [8]).

One of the new discoveries that is beneficial in future emergency and crisis event studies is the authors' quantitative measurement [32] of the proportion of truthful tweets versus number of rumor (unsubstantiated or false) tweets. Out of a sampling of approximately 4k tweets, the authors identified seven true stories and seven false stories, where each story is significant in that it has more than a thousand tweets in the original unsampled dataset. These stories are compared to external reliable sources (ground truth): over 95.5% of the tweets on confirmed events (the set of seven true stories) validate their truth. Conversely, 50% of tweets are observed to refute a story (from the seven false stories) when it is evidently false. These findings suggest that "the Twitter community works like a collaborative filter of information" [32].

The studies by Longueville *et al.* [31] and Mendoza *et al.* [32] complement existing research on crisis events as they help researchers understand the application of microblogging in disaster events, this time from a different cultural perspective (European and South American, instead of North American which was the focus of previous studies) due to the differing usage habits between countries.

Crisis Events: Early Detection and Warning of Disaster There is another aspect of research with regards to Twitter use in emergencies: in early detection and warning of potential emergency situations and to complement existing sensory and surveillance systems. Sakaki *et al.* [33] suggested the use of real-time 'social sensors' in earthquake detection, as Twitter users frequently post details of earthquakes and tremors on Twitter as soon as they feel the event happening. They propose an event detection methodology to validate findings of earthquakes, using Support Vector Machines (SVM) to classify earthquake-related messages with three features: "number of words in a tweet message and the position of the query word within a tweet; the [keywords] in a tweet; [and] the words before and after the query word" [33]. Such processing is done on Japanese language tweets (the scope of this paper focuses on earthquakes, Japan-wide). The authors propose another probabilistic model to handle spatiotemporal information about tweets (using Kalman and particle filters) to predict the earthquakes' trajectory/path, which is then used to create a prototype earthquake reporting system.

Finally, Guy *et al.* [34], researchers for the US Geological Society, developed a "Twitter Earthquake Detector" (TED) to detect earthquakes globally; similar to the research by Sakaki *et al.* [33] as they use both the Twitter user domain and message domain to do so. By listening to the Twitter Streaming API for incoming tweets mentioning earthquake-related terms in several languages, matching tweets are then dumped into a database and sanitized to remove instances of retweets and aggregator users. Similar to previous work such as [35], the authors use the Google Maps API Geocoding Service to find location data in tweets, such as profile location or GPS coordinates. By using a model to find the earthquakes' probable epicenter from the spatial and temporal information found in the collected tweets, and another model dealing with the significance of an earthquake based on user activity, the researchers are able to come up with a report and graphical map overlay detailing findings of a quake. There are several issues identified with the system, however, such as lack of geolocation information in some tweets (which would require a high number of tweet samples to fix), ambiguity of the nouns used to detect quakes, and Twitter activity spikes that occur only after events have elapsed. Despite that, the authors conclude the paper by stating several interesting findings. The Twitter chatter on a quake outperforms traditional sensors in some instances (e.g. a Melbourne quake with a low enough Richter magnitude has been detected in a timely manner, and an earthquake in Indonesia where TED is almost four times faster than traditional detectors). Time series comparison revealed that peaks in Twitter activity correlates to actual quakes and the signal-to-noise ratio of earthquake tweets is high enough to warrant effective detection.

4.2 Pattern Detection and User Clustering on Twitter

More studies have been performed on pattern detection and user clustering based on exhibited demographic and messaging patterns by Twitter users. This section summarizes two of them, one of them newly developed after the introduction of the *Lists* feature on Twitter which allows users to group their friends according to custom per-user categories.

Kim *et al.* [36] have turned Twitter Lists into an invaluable source for detecting commonalities among users in the Twitter community, using information from the user and message domains combined. They posit that lists – a publicly available data source on Twitter – has implicit characteristics which enable us to model commonalities between users as judged by their peers. One unique feature of using lists is that keywords do not have to be explicitly mentioned in tweets; by just looking at a user's association, we can deduce the degree of commonality between him and his peers. For their experiment, they used a dataset of 3.3 million users (approximately 10% of the entire user base of Twitter) belonging to 900k lists; with common list names aggregated together in groups (about 2-3 per group). Chi-square feature selection is applied to the corpus of all tweets belonging to each single user, this is performed on all users belonging to a single list group. To obtain the ground truth, human experimenters associate Twitter users with a particular keyword that best describes the individual user. The list of high chi-square words (i.e. terms with the highest relevance to a list that distinguishes it from the rest) are then compared against the words picked by the human experimenters, which provided an accuracy rate of 0.925. The authors conclude their experiment by stating that "the combination of the Twitter list functionality and the chi-squared feature selection is an efficient tool for inferring user characteristics". Combined with profile information and friend/follow characteristics of the users, this could potentially be beneficial in future research.

Horn [37] contributed to the body of knowledge with his Masters' thesis on automatically classifying tweets – and by extension – users, with potential applications of separating user-generated content from professional content and spam filtering[3]. He proposed two mutually-independent, separate, classification schemes:

[3] *Spam users* on Twitter is defined as users with aggressive following/un-following behavior not commonly found among normal Twitter users; and *spam* on Twitter are links to phishing and malware sites and unsolicited advertisements.

1. *C1* distinguishes between user types, and has the following classification categories: 'news'; 'users' describing everyday users of Twitter; and 'company' containing company promotions, inclusive of spam and sponsored tweets.
2. *C2* distinguishes between pure 'facts' and 'opinions' consisting of subjective user tweets, quotations, or questions.

For his thesis. Horn [37] created a data set of 4,800 tweets from 120 users; and for supervised clustering, chose 'news' users from Twitter accounts of well-known news corporations/newspapers, crawled 'regular users' from several celebrity users as seeds, and highlighted 'company' users from scam websites and spam keywords.

Several findings from analysis in the *C1* category shows that the typical number of distinct words used (and also the rankings of keywords), average length of a tweet, and average interval between tweets are clearly different between samples from all three types of users. The same finding applies to analysis of the two different tweet types (factual vs. opinionated) in *C2*, where factual tweets have almost double the amount of distinct keywords compared to opinion tweets, and that the average time between tweets are almost ten times higher for factual tweets. Different levels of sentiment are also found for each separate category; this will be briefly described in Section 5.2 as it is beyond the scope of this chapter.

Sriram *et al.* [38] has identified eight features (*8F*) of automatically classifying tweet texts to one of the five categories of news, events, opinions, deals, and private messages. Using the 'bag-of-words' concept as a comparison, they applied the Naïve Bayes classification method on the eight features on 5407 tweets from 684 authors; human experimenters provide the ground truth by manually assigning a category to each tweet. The eight features are:

1. author name
2. presence of Internet shortenings, emotions, and slang words
3. presence of time-event keywords (e.g. participant, place, and time information)
4. presence of opinioned words
5. presence of emphasis by capital letters or repeated syllables (e.g. `veeeery`)
6. presence of currency (`$`) or percentage signs (`%`)
7. presence of `@user` notations at the beginning of the message
8. presence of `@user` notations within the message

They conclude their study by stating that their *8F* classification performs better than traditional bag-of-words classification to "classify tweets into general but important categories by using the author information and features within the tweets". We observe that this study [38] mainly deals with the messages on Twitter (with the author information as the only item from the user domain), however findings such as this and in [37] demonstrate the ability of using tweets to indirectly classify users; therefore we encourage the incorporation of features from both users and messages as we believe will increase the accuracy of future research on pattern detection on Twitter.

4.3 Pattern Detection: Applications for Spam Detection and Filtering

A novel area of investigation based on pattern detection in user behavior and their message characteristics is the application on spam detection and filtering. Moh & Murmann [39] and Lee et al. [40] have both performed studies on this, and their findings based on analysis of the friend/follow patterns and live 'honeypot' analysis will be detailed here. Moh & Murmann [39] adapted existing features found from the Twitter users API and synthesized few new attributes/metrics that have been useful in spam detection:

– **Friend/follower count**: derived average friends per day, average followers per day, percent of reciprocal friends
– **Total favorites count**
– **Protected updates flag**
– **Status update count**: derived average updates per day
– **Presence of profile URL**
– **Username**: derived presence of numbers in username string
– Trust metric as a sum of (1/users followed) for all followers of a given user

They use the metrics above not only for individual users, but also for all peers of a given user (a peer is defined as a reciprocal friend, c.f. Kwak *et al.* [8], who has a mutual friend/follow connection). From their findings after applying supervised learning algorithms – JRIP, J48, SOM, and Naïve Bayes as provided by the Weka data-mining package – it is evident that real spammers (as validated by human experimenters) have the following obvious features that discriminate them from the rest, listed by order of decreasing importance:

1. ratio of spammers to legitimate followers
2. average friends per day
3. trust metric (and several weighted variants)
4. friend to follower ratio
5. friends' average friend to follower ratio
6. followers' average of protected users

In a similar vein, Lee *et al.* [40] have identified 14 classes of attributes from both the user and message domain, of which the ones not covered above by Moh & Murmann [39] above are listed:

– **Account age**
– URLs in the **message content**: derived the ratio of URL counts over total tweet count, ratio of unique (non-duplicate) URLs over total tweet count
– `@user` mentions in the **message content**: derived the ratio of `@user` mentions in the last 20 tweets, ratio of unique `@user` mentions in the last 20 tweets
– Average content similarity metric using the standard cosine similarity over the bag-of-words vectors (of tweet content)

Abrol & Khan [41], in the context of their Twitter geocontent study (discussed in Section 4.4), came up with a simple formula to deduce spammers based on the fact that their behavior in the user and message domains are rather atypical of a normal user. As above, the ratio of followers to friends (followees) is rather small; another finding is that a spammer "rarely addresses his messages to some specific people" (using the `@user` or `RT` notations). Their spam confidence level, based on these two findings, is defined as:

$$Z = \frac{1}{\frac{\Sigma\text{followers}}{\Sigma\text{followees}} + \mu\left(\frac{\Sigma\text{reply tweets}}{\Sigma\text{tweets}}\right)}$$

Qualitative observations on spammers and spam users in the Twitter environment have also been performed, both in the honeypot study by Lee *et al.* [40], and also a study on political opinion-spam by Metaxas & Mustafaraj [20]. The honeypots by Lee *et al.* [40] observed that 'social spammers' frequently distribute malware, spam, phishing messages and affiliate programs through contextual 'social'-based spam messages; the following modus operandi is observed:

- duplicated tweets to multiple unsuspecting users
- containing pornographic references in profiles and tweets
- mixing legitimate content with advertising/promotional/phishing spam content
- "infiltrators" behave like a normal user initially but then start disseminating spam after reaching a sizeable following

Meanwhile Metaxas & Mustafaraj [20] observed real-world political opinion-spam disseminated through aggressive campaigns, with a similar modus operandi:

- spammers tend to be unconnected in the friend/follow social graph (compared to other political tweeters who have similarities of user profiles/tweet sentiments with others of the same ideology)
- bogus users relating to the topic being discussed; e.g. spam links with text `coakleysaidit`[4] from 9 spam users sent 929 tweets addressed to 573 unique users in only approximately 2 hours
- spam accounts target users with certain message content

4.4 Viral Information Spread and Memetics

In the earlier literature review [1], related literature not specifically related to Twitter have been studied to learn about the potential nature of viral information spread on the microblogging platform. However, since then, research that deals with the specifics of viral information spread on Twitter via retweeting 'diffusion' and geographic targeting have been authored.

Van Liere [42] studied patterns of information brokering and geographic diffusion of retweets on Twitter by first proposing three patterns: uniform distribution; skewed towards nearby users (local pattern); and skewed towards users

[4] referring to the Coakley–Brown Senate race, January 2010

furthest away from them (information brokerage due to interests which are alike). The author first scans all retweeted messages (by virtue of the text string `RT`), performs sanitization of inaccessible URLs and skewed data, then feeds the conversation topic back into the Search API to retrieve all tweets mentioned. By obtaining user profile information on geographic coordinates for each of the authors, the author determines the exact location and is able to perform a Haversine distance calculation on the exact geographical distance between users. From a dataset of 6,424 geocoded users, he obtained 13.4k retweets stemming from 285 original posts. Temporally, 60% of retweets commence in the first hour after the original tweet was posted, and quickly fades to zero in the time period of about a day. Geographically, he determined that the average retweet distance is approximately 955km, with the median being 1698km, suggesting an 'information broker' pattern (skewed towards users furthest away from them) as discussed earlier.

Van Liere [42] also talked about the motivations behind a retweet, which are three-fold: to vie for attention and increased follower count; to gain influence as a social media filter "who specializes in a particular topic"; and to transfer information from one social network group to another - acting as a bridge between distinct groups [5]. In fact, the first two motivations mentioned can be complemented by one of the findings from Metaxas & Mustafaraj, that "one is much more likely to retweet a message coming from an original sender with whom one agrees (shares political orientation [with])", as the majority of users "were very unlikely to retweet a message that they did not agree with". [20]

Abrol & Khan [41], by observing frequency of tweets mentioning a specific location versus the actual real-world population of the location, was able to come up with a metric called the 'Frequency-Population Ratio (FPR)'. This can then be used as a yardstick to measure any anomalies with respect to patterns of information spread involving issues regarding a particular location; for example the Fort Hood shooting in November 2009 had an abnormal FPR (over 1800 compared to the baseline average of 1), indicating that information (on some event happening in that particular place) has been spreading rather rapidly.

Finally, Stonedahl *et al.* [44] measured viral marketing strategies on four theoretical models alongside the real-world case of Twitter, which would help us understand the viral nature of information spread. The complete details behind the network models used are omitted from this review for clarity, but the findings that set Twitter apart from other theoretical models will be given a discussion. The user network on Twitter is crawled, starting with a random seed Twitter UID between 1 and 10 million; using breadth-first search, they crawled 999 closest nodes from the seed and all 13,343 friendship links between them (reciprocal friends cf. Kwak *et al.* [8]). What Stonedahl *et al.* [44] found out was that Twitter does not take on the form of a random social network; in fact, it has 'hubs' of users with high degree of friends situated close together in the social graph, however 'bridges' between unconnected parts of the network (or "information brokers") are situated rather far away. They conclude that Twit-

[5] Similar to the 'connector' role, c.f. Gladwell's 'Tipping Points' theorem [43]

ter is substantially different to the four networks derived from social network literature, including the regular 'lattice' network and the 'small-world' network.

4.5 Trend and Anomaly Analysis from Twitter and Other Fields

Similar to the previous section, the literature on trend analysis (and even anomaly analysis), hitherto only found in blogs and other social media, is now available with a specific focus on Twitter. One of the earlier works [9] that discussed on trend behavior is complemented and improved upon by several research papers, as will be seen below.

Kumar *et al.* [45] recently performed an experiment to mathematically model the dynamics of conversations (messages and users) in online social networks. One of their findings is that Twitter hashtags exhibit distinctive behavior when they are trending:

- Topics with a high 'preferential attachment' ("messages that have already received many replies [which] are more likely to receive a new reply" [45]) tend to be memes.
- Topics with a high 'copying rate' ("new authors tended to join in often" [45]) have often a stronger 'sense of time', and mainly due with current events.

We now revisit Kwak *et al.* [8], who also performed research on trends and Twitter trending topics, we summarize their findings in detail here. Trending behavior exhibited by topics such as the Apple iPhone launch and the Iran Election controversy are different, despite their similarities in the total number of tweets recorded – the iPhone topic had more users than the Iran Election topic which slowed down in pace. In fact most trending topics have mostly one 'spike' (73% of cases) compared to multiple spikes; about 31% lasted one day or less, but only about 7% persisted for more than ten days. Therefore, the authors came up with a new two-attribute classification scheme to categorize trending topics. *Criticality* is defined in terms of its potential to spread, and *exogeneity/endogeneity* refers to external vs internal "factors that push a topic to the top trending topic list and [cause] the spread of the topic" [8]. Four categories are derived from the two sets of two attributes:

1. **Exogeneous critical**: consists of timely breaking news topics such as about celebrities, causing a single spike but with gradual decay.
2. **Exogeneous subcritical**: consists of 'ephemeral' micromemes and hashtags, usually causing one spike and rapid decaying over time.
3. **Endogeneous critical**: consists of topics of a "more lasting nature [such as] professional sports teams, cities, and brands... [labeled as] persistent news"[8], that have multiple spikes and showing signs of slow decay in activity.
4. **Endogeneous subcritical**: same as endogeneous critical, but the trending period is much shorter.

Huang *et al.* [21] also conducted studies on trends from a hashtag perspective, which reveals more patterns on user activity contributing to tweets over time. By obtaining the standard deviation between timestamps of messages of a single hashtag, they found that small standard deviations of timestamps mark a topic as mostly for conversational purposes (i.e. micromemes, "both adopted and abandoned in a short period of time" [21]) rather than for organizational purposes (i.e. more serious topics such as current affairs). By analyzing the skew of the activity vs. time graph to measure the viral nature of a tag, a negative (left-sided) skew indicates gradual adoption of a tag before reaching its peak of activity, compared to a positive (right-sided) skew which indicates a rapid adoption of a tag before its gradual decline, evident in micromemes. Kurtosis, or the fourth moment, represents the "staying power" of a hashtag [21]; topics with high kurtosis represents bursts of temporal activity (such as trending micromemes which later fade to zero), and a low kurtosis indicates a persistent topic (e.g. the long-discussed H1N1 flu outbreak). As such, the kurtosis metric "can be used to differentiate between micro-memes, recurring tags, or spam." [21]

To detect anomalous users in Twitter and microblogging, related research with dynamics of phone networks by Gupta & Dey [46] can potentially be applied to the user and message domains. They use a set of feature vectors that for both incoming and outgoing communications, identify both the degree of communication and the length of communication (e.g. $f_{OUT}(low, short)$ stands for count of outgoing messages to low-frequency contacts, where messages are short); and a set of global features such as sum of contacts with one-way activity, and total time period. Using a k-nearest neighbor classifier, they successfully applied their methodology on the Enron email dataset and the IEEE VAST 2008 challenge dataset. We posit that this could potentially be applied to Twitter in terms of its message domain (for tweet length) and user domain (friend/follow activity) for similar analyses.

5 Analysis and Modeling on Twitter

Applications of Twitter for analysis and modeling in both the user and message domains have been garnering good reception in the research community; with several applications in regards to existing fields such as opinion/sentiment analysis, information retrieval, and mathematical modelling.

5.1 Recommender Systems and Personalization

Bernstein *et al.* [47] who researched on users' need to customize their user streams found out that more than a thousand tweets can be received and consumed daily if one is an active user. By performing a user study among 78 users, they found out that the average number of tweets daily among this group is around 786 (with a standard deviation of 658), which is too much for a user to consume per day. From their survey, 70 users categorize their Twitter use as personal, while

16 classify them as professional; showing an overlap of users using Twitter for both personal and professional purposes. They have come up with the idea that users will value tweets based on the relevance to their interests, *tie strength* (they would want to see more from users they have a connection with on Twitter), and *serendipity* (interestingly random tweets). Although still a work in progress, they have prototyped a visual personal Twitter feed, which displays 'trends' custom-tailored to an individual user's preferences as a quick 'first-pass' filter to relevant topics of interest.

Phelan *et al.*'s [48] *Buzzer* recommender system is a web-based interface that utilizes Lucene to mine the content of Twitter messages for topics of interest, and allows the user to view personalized streams either in the public timeline, friends' combined messages, or simple term-frequency from RSS feeds. Wu *et al.* [49] has a similar research topic in using the TF-IDF and TextRank mechanisms to perform automated tagging and annotation using keywords extracted and sanitized from a users' collection of Twitter messages.

An observation from this section is that current work such as [48, 49] above primarily operate on the message domain, only incorporating basic ideas from the user domain such as the basic follower network. Interestingly, findings from the message domain such as the tagging and annotation algorithm by Wu *et al.* [49] do allow us to infer the habits and behaviors of users, which would be useful in future studies involving the user domain.

5.2 Opinion Mining and Sentiment Analysis

Traditionally studied in the contexts of blogs, fora, and product reviews (c.f. Pang & Lee [50]), applications of opinion mining and sentimental analysis on microblog content, especially the message domain on Twitter, has been given increasing attention lately. Wasow & Baron's [51] position paper on using tweets as a measure of measuring user sentiment – themed the '*Bruno Effect*' – tries to study the link between microblog sentiment and the real-world 38% decline in revenue for the *Bruno* movie release in 2010. They hypothesize that comments in Twitter which reflect user sentiment relates to the real-world decline in sales performance for that movie. Using the TweetCritics sentimental analysis tool for Twitter, they are able to come up with a model of the difference in revenue (between the first and second days) as a function of the sum of negative tweets recorded in the same time interval. Their secondary analysis reveals that automated sentimental analysis tools such as TweetCritics are comparable to manual sentimental analysis of tweets by human experimenters (in this paper, they are outsourced using the online Amazon Mechanical Turk system). Shamma *et al.* [52] have written a position paper to summarize and supplement their original findings [53] that short tweets can be used effectively to annotate debates based on popular opinion due to the availability of temporal data (i.e. message time-stamp); and that statistics on Twitter messaging activity over a specified time interval can produce cues on the venue, structure, and the activity level in a real-world program, as exhibited in their annotation of the US Presidential Debates [53].

A preliminary short study by Pennacchiotti & Popescu [54] on the usage of sentimental words and "controversial terms (e.g. `trial, apology`) from Wikipedia pages" [54] allowed them to develop an early prototype of a controversy detector for Twitter messages; future work is still ongoing. Bollen *et al.* [55], using *Profile of Mood States* (POMS) analysis to study changes of sentiment on a corpus of 9.6 million tweets (from August–December 2008) have successfully used tweet-based sentiments to model real-life socioeconomic phenomena. Their list of such phenomena is adapted from a list of twenty real-life events, incorporating the corresponding stock market behavior of the Dow Jones Industrial Average and West Texas Intermediate oil price indices. They have found in the study of several events that changes in terms of 'mood states' were evident based on analysis of the sentiments exhibited in public tweets, such as:

1. *US Presidential Election*: doubt before the election (increased mood states of confusion/depression), celebration after the election (increased mood state of vigor, decreased mood state of fatigue).
2. *Thanksgiving holiday*: increased mood state of vigour, decreased mood state of fatigue.
3. *Dow Jones Industrial Average price drop due to bailouts*: increased mood state of depression.
4. *John McCain announcing Sarah Palin as his running mate for the elections*: increased mood state of tension.

To conclude, we observe that all the literature discussed in this section deal entirely with the message domain, as is the traditional practice of using pure textual content for sentimental and opinion analysis. Suggestions for future research include combining such analyses with observations and measurements from the user domain.

5.3 Information Retrieval

Suh *et al.*'s [56] position paper on Twitter with regards to sensemaking has, among other things (such as sentimental analysis discussed earlier), defined the need for search as an information seeking task on Twitter. The key problem in searching microblog messages is the limited length of messages "for search algorithms to work efficiently" [56]. Term expansion is suggested as an approach to handle situations involving Twitter search.

Golovchinsky & Efron [57] have identified similar difficulties in IR on Twitter from a user study, that Twitter search is becoming more frequently used by users trying to locate information such as "[for] events... for people, and for trending topics" [57]. They also identify weaknesses from the Twitter API (as previously discovered in e.g. [9, 8]), one of the important ones include a mere two-week backdated window for search results. Hashtags were identified to improve search experience; the caveat is that the "small size of each tweet makes it hard to estimate the relevance of that tweet" and that the traditional Twitter search interface "[makes] it difficult to extract key memes or documents that characterize an event" [57]. These points are also summarized by researchers such as Bohringer & Gluchowski [58] who identified potential in the area of search and information retrieval; in the authors' words, searching "for a subjectively important information in the data stream" in the "proverbial haystack" [58]

An applied study on information retrieval by Sharifi *et al.* [59] focused on the development of an experimental automatic summarization system that takes trending topic keywords, and applies phrase reinforcement algorithms to automatically generate summaries of tweets discussing a particular topic. As their system is still experiemental, there is no conclusive results as yet.

5.4 User Ranking and Topic Modeling

Ritter *et al.* [60] performed research on unsupervised modeling of conversations in Twitter to detect dialog structure between users. Using 1.3 million conversations from the Twitter public timeline, their first discovery is that each conversation has a range of between 2–243 posts. Then, they sanitize their data by removing non-English tweets, before clustering together misspelled words using the *Jcluster* algorithm. Conversation and topic modelling using LDA and Bayesian methods is performed on the resulting message set. The authors found that coherent patterns do emerge from messages between users on Twitter, and that unsupervised modeling of dialogs on Twitter is indeed possible even with a large dataset. As a result, the authors also come up with a ten-act conversation and topic model corresponding to different features (e.g. reactions versus questions), which will be discussed in further detail in Section 6.4.

Puniyani *et al.*'s [61] research is in a similar vein, in that they construct a LDA topic model over 837k Twitter messages to identify latent themes within. Due to the small size of a Twitter message, their approach aggregates all messages from a user into one 'document', learns the latent topics that characterize authors, and infers the underlying user network structure, user similarity, and connections between users with similar latent topics. Their findings are inconclusive as their experiments are still in progress.

6 Human Factors on Twitter

6.1 Parallels with Other Communication Channels and Social Networks

Communication Patterns This topic is hitherto not discussed in the earlier literature review by Cheong & Lee [1]; since then research on related communication channels and networks have been conducted, which could give us an insight into the similar human factors involved, such as motivations and usage practices. Three of the four papers below in this section deal with cellular phone communication networks, which have been available for a few decades.

Rangaswamy *et al.* [62] studied an Indian SMS service, *SMS Chatter* that parallels the structure of Twitter messages (i.e. short message lengths), but is not an online social network *per se*. The basic premise of this service is that users subscribe to groups or 'chatrooms' and send their messages to subscribers within the same group, not unlike Internet Relay Chat, which is clearly different in structure than Twitter's friend/follow network. From an initial attempt to categorize such groups into one of six categories: news, fun, cricket (sport), computing, educational, and business [62]; the authors found that over time, such discussion categories begin to evolve and that some groups will have overlapping categories due to such evolution. In the authors words, some groups "while originating and evolving around a specific interest area, reach out to include various types of content from other groups as well, some even unrelated to the core interest area, to appeal to a wider audience and prospective members." [62]. They present a couple of case studies illustrating their point:

– *Sartaj* group (200k members): started as a 'fun' group forwarding anecdotes and jokes, has evolved into using advertisements as a potential business model.
– *Vignesh Teacher* group (65k members): started as an 'educational' group disseminating academic information, has evolved into being a prospective advertising tool for promoting jobs.

In closing, the authors conclude that "the entwining of 'fun' and 'business' [referring to the blurred boundaries of discussion group categories] further points to the seamless fusion of a variety of content categories", which increases a group's potential for social networking.

Bentley & Metcalf [63] have performed a field study on three kinds of sharing behavior exhibited by people in terms of mobile communication: motion (sharing one's location), music (sharing information about music currently listened to), and photos. By letting their test particpants use customized Motorola phones to share these types of information, they infer that users directly or indirectly give away certain cues about their current state:

1. **Motion presence**: gives hints to user location, activity, availability, destination (and estimated time of arrival).
2. **Music presence**: gives hints to user location, activity, and availability.
3. **Photo presence**: gives hints to user location, activity, and the presence of people around them.

These findings could be of benefit to microblogging and Twitter research, as Twitter provides (or emulates, with the help of third-party services) some of these 'presence sensors': e.g. location can be exposed by GPS data in tweets, and music/photos can be published using third party services such as TwitPic and Last.fm.

Battestini *et al.* [64] performed a larger-scale study in which the data set comprises of 58.2k SMS messages from 70 participants (age range 17–26 years) over a five-month period. They firstly examined the type of contacts that form the bulk of SMS communication, and found out that all 100% of users send SMS messages to their friends, 60% of users to their family members, while only 24%

of users contact colleagues or workmates. Reasons for SMS chatter includes asking questions or answering them, promoting 'ambient intimacy' with contacts, and chatter about everyday minutiae/trivium. Conversations are categorized in individual categories: with activity planning being the highest proportion (31.7%), followed by relationships, chatting, school/jobs, places, and seeking information (as the top six). However, if we observe simultaneous conversations (i.e. an open-ended message sent to multiple recipients as a start), the categorization scheme becomes different. The following is a non-ranked list of such categories: future plans, sports scores, greetings, thank you messages, big incidents (e.g. robberies), looking for items, announcements, future communication (e.g. broadcasting new phone numbers), and chain letters (spam). Quantitative studies by the authors include tabulating average four-month SMS totals for a person (848.6, $\sigma = 1000.5$); average message length (50.9 characters, $\sigma = 46.2$); and average number of contacts per person (47.1 people, $\sigma = 35.3$).

Again, despite the differences in the network structure for SMS contacts and Twitter friends/followers, the types of messages sent and the comparative statistics for SMS messaging activity could prove beneficial in related microblogging studies.

Adoption Habits A study on Facebook adoption by Bozkir *et al.* [12] reveals the average demographic of a real online social network – Facebook – which could complement findings from Section 3.2 on Twitter's estimated demographics. Several statistics were given as below, on a survey of 570 users:

1. **Gender:** the proportion of males to females is roughly equal
2. **Age bracket:** 74.1% of the surveyd population are in the (18–25) bracket, while 20.53% belong in the (26–35) bracket.
3. **Usage frequency:** 48.77% log in at least once per day, with 25.26% once per week at least.
4. **Period of daily use:** 32.28% for 15 minutes or less on average, 39.82% for 15–30 minutes.
5. **Educational background:** 70.35% are pursuing a Bachelor's degree at the very least, and 23.16% are in their Masters degrees.
6. **Group membership:** surprisingly, 99.82% (all but one user in their sample) is part of a Facebook group.

The authors also hypothesized that several of these statistics have a degree of correlation with others, for example usage time has a correlation with gender/education level, while usage frequency correlates with age/period of daily use; and conclude by suggesting future research on association rules that can model adoption of Facebook as a social network.

6.2 Presence and Outward Sharing

We have surveyed four papers dealing with Twitter which provide an insight into the usage intentions of microbloggers in terms of promoting an online presence and outward sharing of information.

Firstly, Andre's position paper [65] deals with the two topics of "well-being, and status feedback" in terms of microblog messages. In their preliminary study on status feedback using a Web interface which allows users to give ratings (negative, neutral or positive) to tweets by others; they found that the amount of positive feedback outnumber negative ones approximately by a factor of four. They hypothesize that a Twitter user, in their words, "only follow people they are interested in... are friends with many of the people they follow and thus are likely to 'play nice', and people are more comfortable giving positive feedback"[65]. Next, Andre shifted his focus on expression of well-being with tweets, using a custom tagging system. From this experiment, "participants [were] reporting anecdotes of value in self-reflection at the time of update as well as over time" [65], showing that users of microblogs tend to shift their personal state and experience 'online' when engaging in microblogging.

Subramanian & March [66] have performed two studies "to understand what people want to share, with whom, and what challenges they currently face with existing sharing applications" [66]. The first dealt with the researchers 'shadowing' ten test subjects in real-life for a few hours each, capturing photographs and recorded their activities with the subjects' full knowledge. The second study involves nine participants utilizing an iPhone-based photo-sharing application which allows them to take pictures (15–30 daily), and annotate them with a description of what to share and to whom. What they found from these studies is that Twitter usage has evolved from originally "sharing day-to-day intimate details with a small close-knit group... to sharing the mundane and intimate with both close friends and complete strangers" [66]; and that status updates on Twitter are often "carefully crafted" to reflect different aspects or personas of oneself. A list of motivations for outward sharing, which will be detailed in Section 6.5, has been uncovered in this research. Besides these, sharing habits are also influenced by different individual styles and target audiences of their tweets. Other findings from both these studies include the fact that current microblog technologies can't target specific groups of people for different types of updates (contrasting with social networks with custom privacy controls, say Facebook) and the difficulty in defining clear groups from their contacts in the first place; that users tend to 'push' only interesting content to strangers while allowing people close to them to 'pull' more personal or trivial content; and that it is difficult to filtering the level of detail to their different target audiences (e.g. friends versus work superiors).

Ramsden's [67] survey of 17 peoples describing "how and why [they use] Twitter", revealed three main motivations for using a microblog service such as Twitter: to keep other people updated with their goings-on (60% of their sample), directed communications with other users (using public `@user` messages, 22% of their sample), and to publish statuses for personal consumption (e.g. observations/ideas, 7% of their sample). Twitter is also sometimes used in conjunction with (or integrated into the usage of) other services, such as Facebook. An interesting finding relating to [66] above is that 82% of the surveyed users are conscious of "of the different audiences that followed them" [67] and try to avoid publishing e.g. overtly personal statuses for fear of them reaching the wrong audience.

6.3 Privacy Concerns

With the high instances of outward sharing on Twitter, user privacy remains an issue which needs to be paid attention to, as with all types of communication and social networking. Lawler & Molluzzo's [13] study on first-year college student perceptions of privacy on Facebook, MySpace and Twitter, revealed that on the whole, privacy on online social networks are not clearly known nor understood amongst its users. Despite their findings mainly focusing on personally-identifiable information on Facebook, certain pieces of information (such as real name and current location) are available from a microblogging site such as Twitter.

Humphreys *et al.* [6], highlighted the fact that combinations of certain information on Twitter have unexpectedly caused a privacy breach resulting in negative side effects (such as a recent burglary case where the burglar allegedly knew his victim was away due to his Twitter updates). They are of the opinion that users of services such as Twitter fail to inform themselves that privacy settings can be changed and that changes to privacy settings may be deemed too difficult. Humphreys *et al.* highlighted that in August 2009, only less than 8% of the total users have private accounts, a sheer contrast to the time period of January 2007 where 40% of Twitter users are private. As per [13], they realize that temporal information and real names are readily obtained from a public user account. Their experiment focused on manually coding 728 tweets for presence of personally identifiable information (and contact details), locations, names (both nicknames and real names), time (specific nouns or date ranges), and personal pronouns (self-references). Most of their dataset did not contain any personally identifiable information; however they found out that co-occurences of such features, such as personal pronouns with time (15%), personal pronouns with location (10%), and personal pronouns with location and time (3%) would prove significant in ascertaining someone's presence at a given time or place. The authors conclude this paper by postulating that in the worst-case scenario, 3% of overall tweets with mention of both location and time, will extrapolate to approximately 360k potential privacy-risking tweets daily (at the current rate of Twitter activity); and that if data from a user's profile is used together with the tweet, the risk of privacy breaches may be compounded.

'Zooming' and Selective Broadcasting An issue identified in Subramanian & March [66], Ramsden [67], and McAllister [68] is the inability of Twitter and current microblogging services to allow selective filtering of information. As described above, for example, certain users might want to send detailed private tweets for friends and family; while choosing to provide a more vague description for workplace colleagues; and blocking it for everyone else. Subramanian & March writes that this allows "benefit of plausible deniability and allows for the notion of ‘saving face’... and allows the right level of information to be delivered to each audience member” [66]; they proposed the creation of a ‘zoom’ feature that allows the user to customize the level of ‘zoom’ (detail) for each group of contacts they have.

6.4 Social Information Needs and Wants

Two papers reviewed in this section investigate the behavior of Twitter users who use Twitter as an information source; i.e. performing ‘social search’ on Twitter. Kaufman & Chen [69] investigated the role of Twitter “as a tool for capturing comprehensive locality information” [69]; this is achieved by proposing that Twitter can be used to provide location-specific information for a newcomer to successfully better understand the area he is in. Their design proposes the use of a mobile client to help the user in two different methods. “Explore City” is used to reveal places of interest, food, and information about a city based on localized tweets which are arguably better than conventional geographic tools such as Google Maps. “Explore Space” on the other hand focuses on one specific nearby venue of interest, and allows the user to explore detailed information from others’ tweets about that venue. Although still a design concept, this research has shown the effectiveness of Twitter in providing social, location-based, information.

Wilson [70] performed an early “analysis into how people describe and converse about their own information needs” on Twitter. By analyzing about 189k unique tweets from about 163.5k users (with approximately 15k being retweets), Wilson checked for occurrences of, and frequency of “different terms as people describe their searching actions” and how the different terms relate to different search interests. By order of popularity, the top six keywords/terms describing search and types of information sought are as follows:

1. *studying*: exams, scientific studies
2. *hunting*: sport
3. *finding*: ‘finding out’ about things
4. *searching*: food, people, music/pictures, friends
5. *looking*: people, jobs, technology
6. *exploring*: places

Although a work still in progress, Wilson [70] provided us with an insight of the types of information sought on Twitter, as well as the different synonyms frequently associated with social information needs on Twitter.

6.5 Summary of Qualifying Usage Behavior

Westman & Freund [71], in a very detailed study, conducted ‘genre analysis’ on Twitter updates by analyzing six ‘genre facets’ (i.e. the *who, what, when, where, why,* and *how*), manually coding a dataset of 5541 unique tweets. They examined behavior pertaining to the six facets, performed categorization of each tweet for each facet, and summarized their findings in the following set of statistics, quoted verbatim as follows [71]:

1. **Why**: Purpose (mutually exclusive)
   - Information sharing 56 %
   - Conversation 37%
   - Information seeking 8%
2. **What**: Content (mutually exclusive)
   - Personal information or commentary 76%
   - News and public information 18%
   - Business and promotion 4%
   - Issues related to Twitter 3%
3. **Who**: Participant [types] (not mutually exclusive)
   - Popular (more followers than sample median, 106) 50%
   - Experienced (more tweets than sample median, 1129) 50%
4. **When**: Timing (mutually exclusive)
   - Daytime (sent between 8 am and 8 pm) 54%
   - (*conversely, night time: 46%*)
5. **Where**: Place (mutually exclusive)
   - Web (sent from the twitter.com website) 41%
   - (*conversely, other clients: 46%*)
6. **How**: Form (not mutually exclusive)
   - Addressing another Twitter user (indicated by `@`) 46%
   - Linking (URL) 23%
   - Channels [hashtags] (`#`) 9%
   - Retweeting (`RT`) 8%
   - Emoticons (e.g. `:)`) 5%

The authors have found that, if taken into consideration co-occurrences of different facets listed above, they can be summarized in a list of five common genres [71]. This, as far as we know, is (to date) the most comprehensive list of usage genres. We have also performed historical definitions of genres/categories/types of information shared from the earlier literature review [1], and also papers such as [64, 72, 60, 38, 66]. The following list summarizes the findings from Westman & Freund [71], adapted to fit additional definitions from earlier work mentioned; which in our opinion, are absent from the provided genre summary but are evident in Twitter use. Existing definitions that fit into Westman & Freund [71] are included as a subpoint to illustrate their relevance.

1. **Personal Updates** (Westman & Freund [71]): sharing of personal information and also personal opinion, often by more sparsely connected users
   - We redefine this to include opinions and causes which are personal in nature [66].
   - The historical definition of this is "*daily chatter*" (c.f. Java *et al.* [4]).
   - This includes Subramanian & March's [66] definition of "*highlighting a personal moment, providing light entertainment, revealing one's intimate inner thoughts, championing a cause, tracking a professional or personal goal and ... telling people what you are doing*".
   - This includes Ritter *et al.*'s [60] dialogue act of "*status*" .

- This includes Ehrlich & Shami's [72] "*status*" category.
- This includes Sriram *et al.*'s [38] categories of "*Opinions (O)*".
- Examples: `Watching my baby fall asleep` / `I just changed my twitter page bkgornd` [sic] `and now I can't stop looking at it, lol!!`

2. **Directed Dialogue** (Westman & Freund [71]): conversation about personal matters, addressed to certain user(s), and part of a larger tweet stream
   - The historical definition of this is "*conversations*" (c.f. Java *et al.* [4]).
   - This includes Subramanian & March's [66] definition of "*coordinating with someone*" .
   - This includes Battestini *et al.*'s [64] conversational categories of "*planning, relationships, chatting, communication, health, money [and] current status*"; the key difference between this and Personal Updates above is the fact that it belongs to a threaded conversation rather than a singleton [8] by itself.
   - This includes Ritter *et al.*'s [60] dialogue acts of "*question, reaction, comment, answer, response*" .
   - This includes Ehrlich & Shami's [72] "*directed posts*" category.
   - This includes Sriram *et al.*'s [38] categories of "*Private Messages (PM)*".
   - Examples: `@anonymous aww thank you so much!! :)` / `DWL!! what song is that??`.
3. **Real-time Sharing** (Westman & Freund [71]) **and Factoids**: posting of current news, facts, and information (with or without links/hashtags), often originating from custom applications and connected users
   - The historical definitions of this is "*sharing information/URLs & reporting news*" (c.f. Java *et al.* [4]); we have grouped these two into one category as per [71].
   - We redefined this to include retweets (rebroadcasted tweets) and factual tweets (e.g. first-person breaking news, facts and figures) even though hashtags and URLs may not be present.
   - This includes Ritter *et al.*'s [60] dialogue act of "*Reference Broadcast*" .
   - This includes Ehrlich & Shami's [72] "*provide information and retweet*" category.
   - This includes Sriram *et al.*'s [38] categories of "*News (N) & Events (E)*".
   - Example: `#Money Reliance Industries, India's biggest conglomerate http://bit.ly/58jENy`.
4. **Business Broadcasting** (Westman & Freund [71]): posting of business information often via links, channels and retweets, during daytime hours
   - After observing literature such as Horn [37] who segmented "'company' users" with similar characteristics, we posit that such postings have a significant proportion of spam/marketing behavior involved.
   - This includes Subramanian & March's [66] definition of "*crafting one's professional or personal brand*".
   - This includes Sriram *et al.*'s [38] categories of "*Deals (D)*".
   - Example: `6 more hours till our #blackfriday buzzzzz day #clothdiapers sale at www.kellyscloset.com`.

5. **Information Seeking** (Westman & Freund [71]) **and Information Provision** (Subramanian & March [66]): questions and requests for mainly personal information and targeted solutions
   - We expanded this definition to include the "*solutions contributing to a targeted/specialized area*" [66].
   - This includes Battestini *et al.*'s [64] conversational categories of "*information seeking [and knowledge about] school/jobs, places, sports/TV/news, food*".
   - This includes Ritter *et al.*'s [60] dialogue act of "*question to followers*" .
   - This includes Ehrlich & Shami's [72] "*ask question*, [and to a certain extent] *directed question*" category.
   - Examples: `scholarship to lasalle? can a 14 years old girl attend the competition?` / `anyone using google voice?`

To conclude this section, we have merged several classification/categorization schemes into the genre definitions in [71] in an attempt to summarize and group together related schemes to help in further related research. Also, we observe that the 'genre facets' described above [71] – which primarily operate on the message domain of tweets – is of potential use for pattern detection (c.f. Section 4.2) when combined with additional information from the user domain.

7 Practical Applications and Usage of Twitter

7.1 Visualization and Computer-Human Interaction (CHI)

As previously reviewed in [1], applications of Twitter data from both users and messages in terms of visualization, CHI, and electronic art are commonplace. Several more studies have emerged since [1] was published. From our earlier discussion of Bernstein *et al.*'s work [47], we now elaborate on the need for graphical visualization of Twitter feeds. Their proposed *Eddi* system, based on user surveys, allows users to browse tweets aggregated by topics, which are personalized based on their criteria for value (e.g. relevance to interests, connection to other users, or serendipity, as defined above [47]). Figure 1 shows an example of their visualization prototype.

Mathioudakis & Koudas' [73] *TwitterMonitor* is another example of visualization methods on Twitter. Compared to *Eddi*, *TwitterMonitor* does not personalize streams per user, however it reflects the overall trend of the current public timeline of messages. It attempts to group 'bursty' keywords into related groups based on their co-occurrences within the live Twitter message stream, parsing about 10 million messages per day. Such 'bursty'' keyword groups are then analyzed using content extraction algorithms to label them with appropriate keywords that reflect the overall trend by the group. Their live demonstration version at the *SIGMOD'10* conference is illustrated in Figure 2.

Lastly, we'll briefly describe a collection of visualization algorithms by Donath *et al.* [74], the first two of which are directly related to Twitter.

Fig. 1. Bernstein *et al.* with their personalized *Eddi* Twitter interface [47]

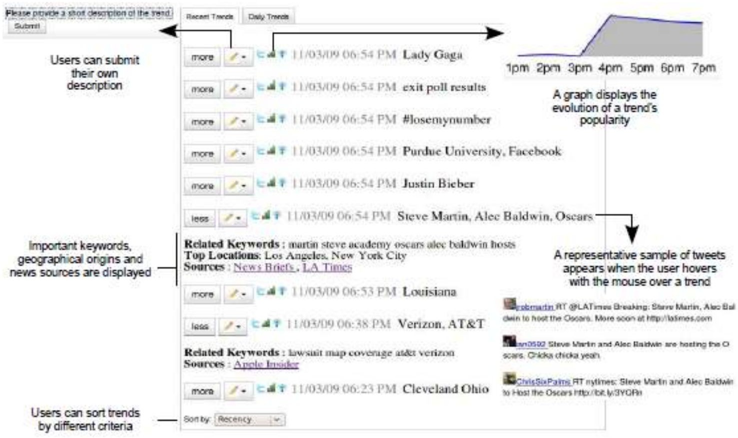

Fig. 2. Mathioudakis & Koudas' [73] *TwitterMonitor* visualization at work.

– *Lexigraphs*: a "group portrait of users of Twitter... shown as a silhouette outlined in words derived from their updates [and] animated by the rhythm of their postings" [74]. This provides the viewer with an opportunity to view the Twitter user from the perspective of his conversation topics on Twitter; the researchers however choose to populate the portrait with "words the subjects use with unusual frequency" [74]. *Lexigraphs* is shown in Figure 3.
– *Mycrocosm*: an introspective visualization of a user's own "everyday 'personal statistics' using simple graphs to display their data" [74] where users can freely choose what kind of information is exhibited on Twitter, and by extension, for *Mycrocosm*. *Mycrocosm* is shown in Figure 4.
– *AuthorLines* and *Themail* shows the timeline of reply and communication habits of users in a discussion forum, and in email, respectively. These 'data-portraits' enable a user to understand his or her own behavior, and allows others in the community to also discern an individuals' behavior in the context of the community.
– *Conversation Maps* are akin to conventional tag clouds that highlight important words in a users' textual communications, with the added twist of having these tag clouds linked to other tag clouds representing other users; in other words it can be seen as a network of tag clouds pertaining to a conversation a user has with his contacts.
– *PeopleGarden* employs word-detection algorithms to determine the emergent features of messages from a particular community (e.g. emotion or political affiliation) to visualize its participants in terms of a 'garden'.

We opine that the latter three can be adapted in future research to deal with microblog data, as Twitter allows the easy exploration of communities (e.g. from the Lists feature or other forms of community detection), and allows the conversations and chatter between users to be made easily available from the user social graph (or from related research e.g. [60]).

7.2 Microblogging in Organizations

Microblogging, which started as a form of intra-organization communication – Twitter was developed as an internal communication tool at Odeo [17] – was recently given research on its characteristics in an organizational context. Thom-Santelli *et al.* [75] studied the IBM *BeeHive* internal lightweight microblogging network from the cross-cultural perspective of three IBM branches: China, India and the United States. The study comprised approximately 60k users (6-13k active at one time) with a total of 150k comments posted. They found out that the microblogging behavior among users differ based on their cultural norms, for example the Indian branch of IBM has users which post informal or more expressive posts; as compared to the ones in the US site. They conclude that "familiarity with the characteristics of other social networking/microblogging sites" [75] influences the internal organizational microblogging behavior of users (e.g. US users are more familiar with Twitter tend to post updates on 'what are you doing?' compared to Indian users' adoption of Orkut which have more

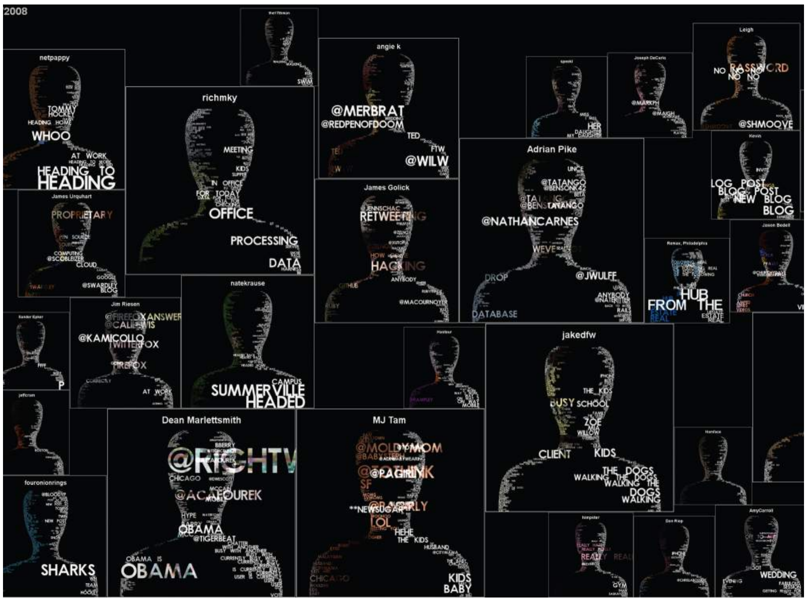

Fig. 3. *Lexigraphs* by Donath *et al.* [74].

Fig. 4. *Mycrocosm* by Donath *et al.* [74].

personally-expressive posts). Cultural power distance also plays a role in influencing the types of status messages created.

For companies which are just starting to adopt microblogging, Zhang *et al.* [76] performed a survey on the adoption of microblogging using the Yammer service in companies. By studying a *Fortune 500* company through a 13-month Yammer data log and some interviews with adopters, they found out that adoption of microblogging in an organization is through four progressive stages, i.e. "initial adoption (registering an account), continued use (continue login for either reading or posting), contributing (posting content; following others), [and finally] promoting (inviting others)" [76]. They also found out that hubs in the user network – "individuals who invite many people to participate" [76] – play a vital role in the adoption of microblogging; however, further study is required to investigate whether these hubs correspond to superiors or high-ranking employees.

Ehrlich *et al.* [72] performed a comparative survey to distinguish between internal organizational microblogs and Twitter by analyzing the contents of over 5k microblog messages by employees of an organization, which are randomly sampled (approximately 3.1k messages from Twitter, and 2.2k originating from an internal microblog site). This dataset of tweets have been authored by 1257 users, of which 25 are interviewed. Prior research [4, 77] has come up with four categories of chatter, conversation, sharing (of information), and news; which are then adapted by Ehrlich *et al.* to create a new list of six categories:

1. Status (i.e. answering 'what are you doing?')
2. Information (comments, opinions, news, and links)
3. Retweets
4. Asking questions
5. Directed tweets (`@user` messages)
6. Directed questions (combination of both directed tweets and questions)

The tweets are manually coded to determine which category suits each individual one; and the time and presence of internal organizational jargon or lingo are studied to provide context for tweets. The results of their study is as follows:

- **Internal/organizational microblogs**: used to chat about company-related subjects, things about work, to get to know colleagues, help colleagues solve problems, and to connect with colleagues (in the context of mobile workers). This form of communication can be summarized as providing information and enabling phatic communications (with less 'noise' and background chatter)
- **Public microblogs such as Twitter**: used to discuss news in real-time, to have conversations/chats with other people, publish personal statuses, directed conversations, and mentioning trivium and information.

7.3 Applied Microblogging in Science, Education, and Governance

This ssection details briefly the role of applied microblogging in science, education (especially the tertiary sector), and governance; as descriptions of such applications in current literature enables us to better understand current microblogging practice by such institutions.

Science Vertesi [78] wrote a position paper on the usage of microblogging by NASA in building a public presence. Examples of Twitter use by NASA include creating Twitter accounts for each of their Mars Rovers to act as an online 'persona'. Several interaction patterns have been discussed by Vertesi [78] in this regard, such as: differentiation of private versus organizational tone of voice for such Twitter accounts, the behavioral patterns of retweets from these accounts, their dissemination of information (retweets and URLs), and the scale of followers for such accounts (tens of thousands, with the highest being about 40k for the *Phoenix Rover*).

Government Wigand [79] has a write-up on the adoption of Twitter by government agencies in the United States. This paper starts with statstics of current Twitter adoption in the USA: almost 20% of Twitter users are from the USA; and that based on *GovTwit*, a directory of governmental 'users' of Twitter, 2,349 users have contributed to more than 192k tweets with about 28 million followers in total. Several examples of US government agencies who have adopted Twitter to reach out to citizens are NASA's Phoenix Rover account (c.f. [78] in the last subsection), the Armed Forces Personnel Administration Agency, and the US State Department which is "one of the main channels to disseminate information about the [Haiti earthquake] emergency" [79]. According to Wigand [79], four major roles of Twitter use in US governmental agencies have been identified: "extending the reach of communication; updating, broadcasting and sharing information through networks; building relationships; [and] collaborating with stakeholders".

Education Du *et al.* [80] has opined that Twitter in the classroom empowers each individual student with the 'right' to say something and express themselves. Ebner *et al.* [81] in a case study on microblogging use in a tertiary setting, found that an average of 7.5 posts per student per working day were generated; of which communication between students and teachers forms a high percentage (60%), and that content which deals with course administration forms 19% of the total. Dunlap & Lowenthal [82] highlighted several use cases of microblogging in an educational environment, such as: asking questions to peers, educators, the community, and even experts (c.f. [67]); facilitating communication (cf. [81]); promoting information sharing, commenting, and dissemination; and 'conference blogging' (i.e. broadcasting live updates from an academic conference). This gives us a summary overview of the increased use of microblogging in educational settings; detailed analysis of Twitter as a facilitator for learning is not covered in this paper as it is beyond the scope of this literature review.

7.4 Potential Fields of Emerging Microblogging Research

Before we conclude this paper, we note down several potential fields of emerging research related to microblogging that have recently been suggested [58]:

- **text mining and semantic analysis on short microblog messages**: the short 140-character length of microblog messages complicates traditional forms of text mining
- **complex event processing**: "each individual micro-blogging posting is in itself constitutes an event" [58] and therefore would be suitable as input for potential complex event processing
- **architecture decentralization and security**: research on this has already started. Xu *et al.* [83] identified centrality as the weakness to a microblogging service, due to: threats to stability such as from denial-of-service attacks, bottlenecks on the system's performance resulting in measures such as 'rate limiting' (which impacts the amount of data that can be collected for research), and single points of failure which can cause the whole microblogging service to fail. Xu *et al.* [83] have designed a working prototype of decentralized microblogging service as a proof-of-concept as a result of their findings.

8 Conclusion

Throughout this literature review paper, we have surveyed more new research on Twitter and microblogging, about twice as many as discussed in Cheong & Lee's [1] earlier review in early 2010. This indicates the massive growth of literature and related research on the subject, as well as applications which are hitherto not researched from the perspective of microblogging.

We have also identified several emerging topics of theoretical research as suggested by several papers in our survey, as well as identified practical applications of Twitter (e.g. in government, education, activism, and for promoting democracy). We expect to see more research related to the surveyed topics, in addition to new and novel applications of microblogging which are currently non-existent, in the near future.

Also, this paper has justified the earlier identification of two separate user and message domains on microblogging services (c.f. [1]). We have also seen the significance of both these domains, their relation to one another, and interdependence, in our evaluation of current literature. Research that merges the study of both these domains are still lacking, in our opinion; which we expect to change in the next iteration of research on microblogging.

References

1. Cheong, M., Lee, V.: Dissecting Twitter: A Review on Current Microblogging Research and Lessons from Related Fields. In: From Sociology to Computing in Social Networks: Theory, Foundations and Applications. Volume 1 of LNSN. Springer-Verlag (2010) 343 – 362

2. Cormode, G., Krishnamurthy, B., Willinger, W.: A Manifesto for Modeling and Measurement in Social Media. First Monday **15**(9) (2010) Available from http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/viewArticle/3072/2601.
3. Huberman, B., Romero, D., Wu, F.: Social networks that matter: Twitter under the microscope. Technical report, Social Computing Laboratory, HP Labs (2008) Available from http://ssrn.com/abstract=1313405.
4. Java, A., Song, X., Finin, T., Tsen, B.: Why we Twitter: An analysis of a microblogging community. In: Proc. 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis, Springer-Verlag (2009) 118–138
5. Krishnamurthy, B., Gill, P., Arlitt, M.: A few chirps about Twitter. In: Proc. WOSN'08. (2008) 19–24
6. Humphreys, L.: Historicizing Microblogging. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
7. Petrovic, S., Osborne, M., Lavrenko, V.: The Edinburgh Twitter Corpus. In: Proc. NAACL HLT 2010 Workshop on Computational Linguistics in a World of Social Media. (2010) 25–26
8. Kwak, H., Lee, C., Park, H., Moon, S.: What is Twitter, a Social Network or a News Media? In: Proc. WWW 2010. (2010)
9. Cheong, M., Lee, V.: Integrating web-based intelligence retrieval and decision-making from the Twitter Trends knowledge base. In: Proc. CIKM 2009 Co-Located Workshops: SWSM 2009. (2009) 1–8
10. Yoshida, M., Inui, T., Yamamoto, M.: Analysis of Tweets Including URLs on Twitter (translated). In: Proc. DEIM Forum 2010. (2010)
11. Pear Analytics: Twitter Study August 2009. Technical report, Pear Analytics (2009) Available from http://www.pearanalytics.com/blog/wp-content/uploads/2010/05/Twitter-Study-August-2009.pdf.
12. Bozkir, A.S., Mazman, S.G., Sezer, E.A.: Identification of User Patterns in Social Networks by Data Mining Techniques: Facebook Case. In: IMCW 2010. Number 96 in CCIS (Communications in computer and information science) (2010) 145153
13. Lawler, J.P., Molluzzo, J.C.: A Survey of First-year College Student Perceptions of Privacy in Social Networking. Journal of Computing Sciences in Colleges **26**(3) (2011)
14. Krishnamurthy, B.: A Measure of Online Social Networks (Invited paper). In: Proc. COMSNETS'09. (2009)
15. Honeycutt, C., Herring, S.: Beyond microblogging: Conversation and collaboration via Twitter. In: Proc. 42nd Hawaii International Conference on System Sciences. (2009) 1–10
16. boyd, d., Golder, S., Lotan, G.: Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In: Proc. HICCS-43. (2010)
17. O'Reilly, T., Milstein, S.: The Twitter Book. O'Reilly Media, Inc., Sebastopol, CA (2009)
18. Baumer, E., Leis, A.: Minimalists and Zealots: Genres of Participation in Following on Twitter. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
19. Heil, B., Piskorski, M.: New Twitter Research: Men Follow Men and Nobody Tweets. Available from http://blogs.harvardbusiness.org/cs/2009/06/new_twitter_research_men_follo.html (June 2009)
20. Metaxas, P.T., Mustafaraj, E.: From Obscurity to Prominence in Minutes: Political Speech and RealTime Search. In: Proc. Web Science Conf. (2010)
21. Huang, J., Thornton, K.M., Efthimiadis, E.N.: Conversational Tagging in Twitter. In: Proc. HT10. (2010)

22. Hughes, A., Palen, L.: Twitter adoption and use in mass convergence and emergency events. In: Proc. 6th International ISCRAM Conference. (2009)
23. Starbird, K., Palen, L., Hughes, A., Vieweg, S.: Chatter on The Red: What hazards threat reveals about the social life of microblogged information. In: Proc. CSCW 2010. (2010)
24. Ems, L.: Twitter Use in Iranian, Moldovan and G-20 Summit Protests Presents New Challenges for Governments. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
25. Lin, Y.R., Tolentino, L., Kelliher, A.: Tweeting Globally, Acting Locally: Booming and Sustaining Disability Awareness through Twitter. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
26. Burns, A., Eltham, B.: Twitter Free Iran: An Evaluation of Twitter's Role in Public Diplomacy and Information Operations in Iran's 2009 Election Crisis. In: Communications Policy & Research Forum 2009, University of Technology, Sydney (2009)
27. Jungherr, A.: Twitter in Politics: Lessons Learned during the German Superwahljahr 2009. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
28. Serbanuta, C., Chao, T., Takazawa, A.: Save the Tweets So You Can Understand the Birds. In: Proc. iConference 2010. (2010)
29. Vieweg, S., Starbird, K.: Microblogging in Mass Emergency. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
30. Kireyev, K., Palen, L., Anderson, K.: Applications of Topics Models to Analysis of Disaster-Related Twitter Data. In: Proc. NIPS 2009 Workshop on Applications for Topic Models: Text and Beyond. (2009)
31. Longueville, B.D., Smith, R.S., Luraschi, G.: OMG, from here, I can see the flames!: A Use Case of Mining Location Based Social Networks to Acquire Spatiotemporal Data on Forest Fires. In: Proc. LBSN '09: Proceedings of the 2009 International Workshop on Location Based Social Networks. (2009) 73–80
32. Mendoza, M., Poblete, B., Castillo, C.: Twitter Under Crisis: Can We Trust What We RT? In: Proc. 1st Workshop on Social Media Analytics (SOMA '10). (2010)
33. Sakaki, T., Okazaki, M., Matsuo, Y.: Earthquake Shakes Twitter users: Real-time Event Detection by Social Sensors. In: Proc. WWW '10. (2010) 851–860
34. Guy, M., Earle, P., Ostrum, C., Gruchalla, K., Horvath, S.: Integration and Dissemination of Citizen Reported and Seismically Derived Earthquake Information via Social Network Technologies. In: Advances in Intelligent Data Analysis IX. Springer-Verlag (2010) 42–53
35. Cheong, M., Lee, V.: Twitmographics: Learning the Emergent Properties of the Twitter Community. In: From Sociology to Computing in Social Networks: Theory, Foundations and Applications. Volume 1 of LNSN. Springer-Verlag (2010) 323–342
36. Kim, D., Jo, Y., Moon, I.C., Oh, A.: Analysis of Twitter Lists as a Potential Source for Discovering Latent Characteristics of Users. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
37. Horn, C.: Analysis and Classification of Twitter Messages. Master's thesis, Graz University of Technology (2010)
38. Sriram, B., Fuhry, D., Demir, E., Ferhatosmanoglu, H., Demirbas, M.: Short Text Classification in Twitter to Improve Information Filtering. In: Proc. SIGIR '10. (2010) 841–842
39. Moh, T.S., Murmann, A.J.: Can You Judge a Man by His Friends? - Enhancing Spammer Detection on the Twitter Microblogging Platform Using Friends and Followers. In: Information Systems, Technology and Management. Volume 54 of

Communications in Computer and Information Science., Springer Berlin Heidelberg (2010) 210–220
40. Lee, K., Caverlee, J., Webb, S.: Uncovering Social Spammers: Social Honeypots + Machine Learning. In: Proc. SIGIR'10. (2010) 435–442
41. Abrol, S., Khan, L.: TWinner: Understanding News Queries with Geo-content using Twitter. In: Proc. GIR'10. (2010)
42. van Liere, D.: How Far Does a Tweet Travel? Information Brokers in the Twitterverse. In: Proc. MSM10. (2010)
43. Gladwell, M.: The Tipping Point: How Little Things Can Make a Big Difference. Back Bay Books, New York, NY (2002)
44. Stonedahl, F., Rand, W., Wilensky, U.: Evolving Viral Marketing Strategies. In: Proc. GECCO '10. (2010) 1195–1202
45. Kumar, R., Mahdian, M., McGlohon, M.: Dynamics of Conversations. In: Proc. KDD10. (2010)
46. Gupta, N., Dey, L.: Detection and Characterization of Anomalous Entities in Social Communication Networks. In: Proc. ICPR'10. Volume 0. (2010) 738–741
47. Bernstein, M., Kairam, S., Suh, B., Hong, L., Chi, E.H.: A Torrent of Tweets: Managing Information Overload in Online Social Streams. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
48. Phelan, O., McCarthy, K., Smyth, B.: Using Twitter to Recommend Real-Time Topical News. In: Proc. RecSys09. (2009) 385–388
49. Wu, W., Zhang, B., Ostendorf, M.: Automatic Generation of Personalized Annotation Tags for Twitter Users. In: Proc. HLT '10. (2010) 689–692
50. Pang, B., Lee, L.: Opinion Mining and Sentiment Analysis. Volume 2 of Foundation and Trends in Information Retrieval. now Publishers Inc., Hanover, Boston MA (2008)
51. Wasow, O., Baron, A.: Can Tweets Kill a Movie? An Empirical Evaluation of the Bruno Effect. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
52. Shamma, D.A., Kennedy, L., Churchill, E.F.: Media, Conversations and Shadows. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
53. Shamma, D.A., Kennedy, L., Churchil, E.F.: Tweet the debates: Understanding community annotation of uncollected sources. In: Proc. ACM Multimedia 2009. (2009)
54. Pennacchiotti, M., Popescu, A.M.: Detecting Controversies in Twitter: A First Study. In: Proc. NAACL HLT 2010 Workshop on Computational Linguistics in a World of Social Media. (2010) 31–32
55. Bollen, J., Pepe, A., Mao, H.: Modeling Public Mood and Emotion: Twitter Sentiment and Socio-economic Phenomena. Technical report, School of Informatics and Computing, Indiana University - Bloomington (2009)
56. Suh, B., Hong, L., Convertino, G., Chi, E.H., Bernstein, M.: Sensemaking with Tweeting. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
57. Golovchinsky, G., Efron, M.: Making Sense of Twitter Search. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
58. Bhringer, M., Gluchowski, P.: Aktuelle Schlagwrter: Microblogging. Informatik Spektrum **32**(6) (2009) 506–510
59. Sharifi, B., Hutton, M.A., Kalita, J.: Summarizing Microblogs Automatically. In: Proc. Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the ACL. (2010) 685–688
60. Ritter, A., Cherry, C., Dolan, B.: Unsupervised Modeling of Twitter Conversations. In: Proc. Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the ACL. (2010) 172–180

61. Puniyani, K., Eisenstein, J., Cohen, S., Xing, E.P.: Social Links from Latent Topics in Microblogs. In: Proc. NAACL HLT 2010 Workshop on Computational Linguistics in a World of Social Media. (2010) 19–20
62. Rangaswamy, N., Jiwani, S., Chowdhury, I.R.: Micro-blogging or GupShup (chatter): Mobile Chattering in India. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
63. Bentley, F., Metcalf, C.: The Use of Mobile Social Presence. IEEE Pervasive Computing **8**(4) (2009) 35–41
64. Battestini, A., Setlur, V., Sohn, T.: A Large Scale Study of Text Messaging Use. In: Proc. MobileHCI10. (2010) 229–238
65. Andr, P., schraefel, m., Dix, A., White, R., Bernstein, M., Luther, K.: Designing for Schadenfreude (or, How to Express Well-being and See if you're Boring People). In: Proc. CHI 2010 Workshop on Microblogging. (2010)
66. Subramanian, S., March, W.: Sharing Presence: Can and Should Your Tweets be Automated? In: Proc. CHI 2010 Workshop on Microblogging. (2010)
67. Ramsden, A.: How and Why are People Using Twitter: A Small Group Study. Technical report, University of Bath (2008)
68. McAllister, B.: Why "The Conversation" isn't Necessarily a Conversation. interactions **17**(5) (2010) 19–21
69. Kaufman, S.J., Chen, J.: Where We Twitter. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
70. Wilson, D.W.: Monitoring technology trends with podcasts, RSS and Twitter. Library Hi Tech News **25**(10) (2008) 8–12
71. Westman, S., Freund, L.: Information Interaction in 140 Characters or Less: Genres on Twitter. In: Proc. IIiX 2010. (2010)
72. Ehrlich, K., Shami, N.: Microblogging Inside and Outside the Workplace. In: Proc. ICWSM 2010. (2010)
73. Mathioudakis, M., Koudas, N.: TwitterMonitor: Trend Detection over the Twitter Stream. In: Proc. SIGMOD10. (2010)
74. Donath, J., Dragulescu, A., Zinman, A., Viégas, F., Xiong, R.: Data Portraits. In: ACM SIGGRAPH 2010 Art Gallery. SIGGRAPH '10 (2010) 375–383
75. Thom-Santelli, J., DiMicco, J.M., Millen, D.R.: Cross-Cultural Analysis of Status Messages within IBM. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
76. Zhang, J., Qu, Y., Hansen, D.: Modeling User Acceptance of Internal Microblogging at Work. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
77. Zhao, D., Rosson, M.: How and why people Twitter: the role that micro-blogging plays in informal communication at work. In: Proc. GROUP'04. (2009) 243–252
78. Vertesi, J.: Tweeting Spacecraft. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
79. Wigand, F.D.L.: Twitter Takes Wing in Government: Diffusion, Roles, and Management. In: Proc. 11th Annual International Digital Government Research Conference on Public Administration Online: Challenges and Opportunities, dg.o '10. (2010) 66–71
80. Du, H., Carroll, J., Rosson, M.B.: Public Micro-blogging in Classrooms: Towards an Active Learning Environment. In: Proc. CHI 2010 Workshop on Microblogging. (2010)
81. Ebner, M., Lienhardt, C., Rohs, M., Meyer, I.: Microblogs in Higher Education - A Chance to Facilitate Informal and Process-oriented Learning? Computers & Education **55**(1) (2010) 92 – 100
82. Dunlap, J.C., Lowenthal, P.R.: Tweeting the Night Away: Using Twitter to Enhance Social Presence. Journal of Information Systems Education **20**(2) (2009)

83. Xu, K., Farkas, D.K.: Blogging as a rhetorical act. In: Proc. 73rd Association for Business Communication Annual Convention. (2008)