

Twitter in Academic Conferences: Usage, Networking and Participation over Time

Xidao Wen
University of Pittsburgh
Pittsburgh, USA
xiw55@pitt.edu

Christoph Trattner
Know-Center
Graz, Austria
ctrattner@know-center.at

Yu-Ru Lin
University of Pittsburgh
Pittsburgh, USA
yurulin@pitt.edu

Denis Parra
PUC Chile
Santiago, Chile
dparra@ing.puc.cl

ABSTRACT

Twitter is often referred to as a backchannel for conferences. While the main conference takes place in a physical setting, attendees and virtual attendees socialize, introduce new ideas or broadcast information by microblogging on Twitter. In this paper we analyze the scholars' Twitter use in 16 Computer Science conferences over a timespan of five years. Our primary finding is that over the years there are increasing differences with respect to conversation use and information use in Twitter. We studied the interaction network between users to understand whether assumptions about the structure of the conversations hold over time and between different types of interactions, such as retweets, replies, and mentions. While 'people come and people go,' we want to understand what keeps people staying engaged with the conference on Twitter. By casting the problem as a classification task, we find different factors that contribute to the continuing participation of users to the online Twitter conference activity. These results have implications for research communities to implement strategies for continuous and active participation among members.

Categories and Subject Descriptors: H.2.8 [Database Management]: Database Applications—*Data mining*

Keywords: Twitter; academic conferences; usage; interactions; retention

1. INTRODUCTION

Twitter, as one of the most popular microblogging services, has been raised as the backchannel of academic conferences [19]. There is a considerable amount of research into understanding the users' behavior on Twitter during academic events. Researchers look into why people tweet [17, 19], and what people tweet from a small number of conferences [5, 12, 23].

In this work, we collect data about a larger set of academic conferences over five consecutive years, and provide in-depth analysis on this temporal dataset. We are particularly interested in the following research questions:

RQ1: Do users use Twitter more for socializing with peers or for information sharing during conferences? How has such use of Twitter during conferences changed over the years?

RQ2: What are the structures of conversation and information sharing networks in individual conferences? Have these network structures changed over time?

RQ3: Do users participate on Twitter for the same conference over consecutive years? To what extent can we predict users' future conference participation?

To answer these questions, we crawled a dataset that consists of the tweets from 16 Computer Science conferences from 2009 to 2013. We examined Twitter in conferences by characterizing its use through retweets, replies, etc. We studied the structure of conversation and information-sharing by deriving two networks from the dataset, conversation and retweet network. Furthermore, to understand the factors that drive users' continuous participation, we propose a prediction framework with usage and network metrics.

As a result of our analyses, we found: (i) an increasing trend of informational usage (URLs and retweets) compared to the stable pattern of conversational usage (replies and mentions) of conferences on Twitter over time; (ii) the conversation network is more fragmented than the information network, and the former becomes more fragmented over the time; and (iii) that the number of timeline tweets, users' centrality in information networks, and number of contacts in conversation networks are the most relevant features to predict users' continuous participation. These results summarize the online user participation of a real research community, which in turn helps to understand how it is perceived in online social networks and whether it is able to attract recurring attention over time.

The rest of the paper is structured as follows: The next section surveys Twitter used as a backchannel in conferences. Then, Section 3 describes the dataset in this study and how we obtained it. Section 4 presents the experiment setup, followed by section 5 which provides the results. Section 6 summarizes our findings and concludes our paper with discussion of the future work.

2. RELATED WORK

Twitter usage has been studied in events as diverse as politics [10, 21], sports [9, 25], and natural disasters [1, 15]. However, research that studies Twitter as a backchannel in academic conferences is closer to our work. Ebner et al. [5] studied tweets posted during the ED-MEDIA 2008 conference, and they argued that micro-blogging can enhance participation in the usual conference setting. They conducted a survey over 41 people who used Twitter during conferences, finding that people who actively partic-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HT'14, September 1–4, 2014, Santiago, Chile.

Copyright 2014 ACM 978-1-4503-2954-5/14/09...\$15.00.

<http://dx.doi.org/10.1145/2631775.2631826>.

ipate via Twitter are not only physical conference attendants, and that the reasons people participate are sharing resources, communicating with peers, and establishing an online presence [17].

Considering another area, Letierce et al. [12] studied Twitter usage by the semantic web community during the conference ISWC 2009. They further analyzed three conferences [11] and they found that analyzing Twitter activity during the conference helps to summarize the event (by categorizing hashtags and URLs shared), and that the way people share information online is affected by the use of Twitter. In another study, Ross et al. [19] investigated the use of Twitter as a backchannel within the community of digital humanists. By studying three conferences in 2009 they found that the micro-blogging activity during the conference is not a single conversation but rather multiple monologues with few dialogues between users and that the use of Twitter expands the participation of members of the research community. With respect to applications, Sopan et al. [20] created a tool that provides real-time visualization and analysis for backchannel conversations in online conferences.

In [24], we investigated the use of Twitter in three conferences in 2012 related to user modeling research communities. We classified Twitter users into groups and we found that the most senior members of the research community tend to communicate with other senior members, and newcomers (usually masters or first year PhD students) receive little attention from other groups, challenging Reinhardt’s assumption [17] about Twitter being an ideal tool to include newcomers in an established learning community.

Compared to previous research and to the best of our knowledge, this article is the first one studying a larger sample of conferences (16 in total) over a period of five years (2009-2013). This dataset allows us to generalize our results to Information and Computer Science and also to analyze trends of Twitter usage over time.

3. DATASET

Previous studies on analyzing the tweets during conferences examined a small number of conferences [4, 12]. For each conference, they collected the tweets that contain the conference official hashtag in its text, for example, *#kidneywk11*, *#iswc2009*. They produced insights of how users employ Twitter during the conference, but their results are limited considering that they analyzed at most three conferences. On the other hand, we are interested in studying trends of the usage and the structure over time, where we aimed to collect a dataset of tweets from a larger set of conferences over several years. Following the list of conferences in Computer Science listed in *cscnf.net*, we used the Topsy API¹ to crawl tweets by searching for the conference hashtag. With regard to the time period of the conference, as summarized by [17], Twitter activities can happen at different stages of a conference: before a conference, during a conference, and after a conference. For this purpose we set the search time window to be seven days before the conference until seven days after the conference, in order to capture most of the Twitter activities about the conference.

Conference dataset. For this study, we focused on the conferences that had Twitter activity from 2009 to 2013. The crawling process took two weeks in December 2013. We aggregated 109,076 tweets from 16 conferences over the last five years.

User-Timeline dataset. We acknowledge that users would also interact with other users without the conference hashtag, and therefore we additionally constructed the timeline dataset by crawling the timeline tweets of those users who participated in the conference during the same period (users from the conference dataset). Table 1 shows the statistics of our dataset. In addition, we publish the detailed information about each conference².

	2009	2010	2011	2012	2013
#Unique Users	1,114	2,970	3,022	5,590	5,085
#Conference Tweets	8,125	18,107	19,308	34,424	27,549
#Timeline Tweets	228,923	608,308	589,084	1,025,259	939,760

Table 1: Properties of the dataset collected in each year.

Random dataset. Any pattern observed would be barely relevant unless we compare with a baseline, because the change might be a byproduct of Twitter usage trends overall. Hence, we show the conference tweets trend in comparison with a random sampled dataset. Several sampling methods about data collection from Twitter have been discussed in [18]. Unfortunately, none of those approaches are applicable to this study, as most of them drew the sample from the tweets posted during limited time period through Twitter APIs. Sampling from the historical tweets (especially for the tweets in the five year period) via Twitter APIs seems to be a dead end. To overcome this issue, we again used Topsy API, for which it claims to have full access to all historical tweets. As Topsy does not provide direct sampling APIs, we then wrote a script to construct a sample from all the tweets in the Topsy archive and to ensure the sampling process is random and the sample acquired is representative of the tweets. To eliminate the diurnal effect on Twitter, we randomly picked two-thirds of all the hours in each year, randomly picked two minutes from the each hour as our search time interval, and randomly picked the page number in the returned search result. The query aims to search for the tweets that contain any one of the alphabetical characters (from a to z). The crawling process took two days in December, 2013. As each query returned us 100 tweets, we were able to collect 5,784,649 tweets from 2009 to 2013. Our strategy was designed to create a quasi-random sample using the Topsy search APIs. To examine the quality of our sample, we compared our dataset with the statistics of Boyd et al. [3]. In 2009, 21.8% of the tweets in our random dataset contain at least one URL, close to 22% as reported in their paper. The proportion of tweets with at least one ‘@user’ in its text is 37%, close to 36% in Boyd’s data. Moreover, 88% of the tweets with ‘@user’ begin with ‘@user’ in our sampled data, comparable to 86% of Boyd’s. These close distributions support the representativeness of our dataset during 2009, then we extended the sampling method for the ongoing years.

4. METHODOLOGY

In this section we describe our experimental methodology, i.e., the metrics used, analyses and experiments conducted to answer the research questions.

4.1 Analyzing the Usage

We examined the use of Twitter during conferences by defining the measurements from two aspects: information usage and conversation usage. We want to use these measures to understand different usage dimensions and whether they have changed over time.

Conversation usage. With respect to using Twitter as a medium for conversations, we defined features based on two types of interactions between users: Reply and Mention ratios. For instance, @Alice can reply to @Bob, by typing ‘@Bob’ at the beginning of a tweet, and this is recognized as a reply from @Alice. @Alice can also type @Bob in any part of her tweet except at the beginning, and this is regarded as a mention tweet. We computed the Reply Ratio to measure the proportion of tweets categorized as replies and the Mention Ratio respectively, but considering mentions.

Information usage. For the informational aspect of Twitter use during conferences, we computed two features to measure how it changed over the years: URL Ratio and Retweet Ratio. Intuitively, most of the URLs shared on Twitter during conference time are

¹<http://topsy.com>

²<https://github.com/xidaow/twitter-academic>

linked to additional materials such as presentation slides, publication links, etc. We calculated the URL Ratio of the conference to measure which proportion of tweets are aimed at introducing information to Twitter. The URL Ratio is simply the number of the tweets with ‘http:’ over the total number of the tweets in the conference. The second ratio we used to measure informational aspects is Retweet Ratio, as the retweet plays an important role in disseminating the information within and outside the conference. We then calculated the Retweet Ratio to measure the proportion of tweets being shared in the conference. To identify the retweets, we followed a fairly common practice [3], and used the following keywords in the queries: ‘RT @’, ‘retweet @’, ‘retweeting @’, ‘MT @’, ‘rt @’, ‘thx @’.

We computed the aforementioned measures from the tweets in the conference dataset (tweets that have the conference hashtag in the text, as explained in Section 3). Following the same approach, we computed the same measures from the tweets in the random dataset, as we wanted to understand if the observations in conferences differ from general usage on Twitter.

4.2 Analyzing the Networks

To answer the research question *RQ2*, we conducted a network analysis following Lin et al. [13], who constructed networks from different types of communications: hashtags, mentions, replies, and retweets; and used their network properties to model communication patterns on Twitter. We followed their approach and focused on two networks derived from our dataset: conversation network, and retweet network. We defined them as follows:

- Conversation network: We built the user-user network of conversations for *every conference each year*. This network models the conversational interactions (replies and mentions) between pairs of users. Nodes are the users in one conference and one edge between two users indicates they have at least one conversational interaction during the conference.
- Retweet network: We derived the user-user network of retweets for *each conference each year*, in which a node represents one user and a directed link from one node to another means the source node has retweeted the targeted one.

The motivation for investigating the first two networks comes from the study of Ross et al. [19], who stated that: a) the conference activity on Twitter is constituted by multiple scattered dialogues rather than a single distributed conversation, and b) many users’ intention is to jot down notes and establish an online presence, which might not be regarded as an online *conversation*. This assumption is also held by Ebner et al. [17]. To assess if the validity of these assumptions holds over time, we conducted statistical tests over network features, including the number of nodes, the number of edges, density, diameter, the number of weakly connected components, and clustering coefficient of the network [22].

We constructed both conversation and retweet networks from the users’ timeline tweets in addition to the conference tweets, as we suspect that many interactions might happen between the users without using the conference hashtag. Therefore, these two datasets combined would give us a complete and more comprehensive dataset. Furthermore, we filtered out the users targeted by messages in both networks who are not in the corresponding conference dataset to assure these two networks only capture the interaction activities between conference Twitter users.

4.3 Analyzing Continuous Participation

To understand which users’ factors drive their continuing participation in the conference on Twitter, we trained a binary classifier with some features induced from users’ Twitter usage and their network metrics. From our own experience of attending the same conference repeatedly, we know that one reason is that we had valuable experience in the past –quality research and social connections, to

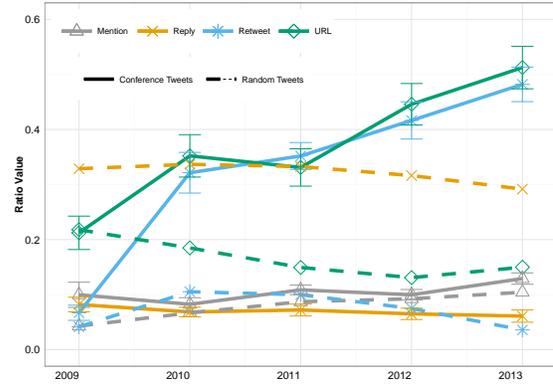


Figure 1: Usage pattern over the years in terms of proportion of each category of tweets. Continuous lines represent conference tweets, dashed lines a random dataset from Twitter.

be more specific. We expect that a similar effect exists with respect to the continuous participation on Twitter. Users’ decision of whether or not to return to the conference possibly depends on their experience with the conference via Twitter in the past – valuable information and meaningful conversations. To capture both ends of a user’s Twitter experience, we computed the usage measures, as described in Section 4.1, and user’s network position [13] in each of the networks: conversation network and retweet network, as discussed in Section 4.2. Measures for user’s network position are calculated to represent the user’s relative importance within the network, including degree, in-degree, out-degree, HIT hub score [8], HIT authority score [8], PageRank score [16], eigenvector centrality score [2], closeness centrality [22], betweenness centrality [22], and clustering coefficient [22].

Dataset. We identified 14,456 unique user-conference participations from 2009 to 2012 in our dataset. We then defined a positive continuing participation if one user showed up again in the same conference he or she participated in via Twitter last year, while a negative positive continuing participation if the user failed to. For example, @Alice posted a tweet with ‘#cscw2010’ during the CSCW conference in 2010, we counted it as one positive continuing participation if @Alice posted a tweet with ‘#cscw2011’ during the CSCW conference in 2011. By checking these users conference participation records via Twitter in the following year (2010-2013), we identified 2,749 positive continuing participations. We then constructed a dataset with 2,749 positive continuing participations and 2,749 negative continuing participations (random sampling [7]). **Features.** In the prediction dataset, each instance consisted of a set of features that describe the user’s information *in one conference in one year* from different aspects, and the responsive variable was a binary indicator of whether the user came back in the following year. We formally defined the key aspects of one user’s features discussed above, in the following:

- Baseline: This set only includes the number of timeline tweets and the number of tweets with the conference hashtag, as the users’ continuous participation might be correlated with their frequencies of writing tweets. We consider this as the primary information about a user and this information will be included in the rest of the feature sets.
- Conversation: We built this set by including the conversation usage measures (Mention Ratio and Reply Ratio) and network position of the user in the conversation network.
- Information: Different from the previous one, this set captures the information oriented features, including the information usages (Retweet Ratio, URL Ratio) and user’s position in the retweet network.

Feature	2009	2010	2011	2012	2013	
Conversation	#Nodes	165.313 ± 50.358	323.688 ± 100.481	385.625 ± 100.294	649.313 ± 202.518	622.188 ± 142.485
	#Edges	342.188 ± 126.758	660.625 ± 240.704	688.500 ± 227.768	1469.000 ± 643.431	1157.813 ± 344.484
	In/Out degree	1.446 ± 0.153	1.567 ± 0.161	1.502 ± 0.086	1.646 ± 0.144	1.618 ± 0.088
	Density	0.044 ± 0.020	0.010 ± 0.002	0.007 ± 0.001	0.005 ± 0.001	0.004 ± 0.000
	Clustering Coefficient	0.066 ± 0.015	0.086 ± 0.014	0.074 ± 0.008	0.070 ± 0.009	0.078 ± 0.006
	Reciprocity	0.172 ± 0.034	0.210 ± 0.029	0.237 ± 0.023	0.195 ± 0.022	0.203 ± 0.017
	#WCCs	4.750 ± 0.911	13.438 ± 3.243	16.625 ± 4.118	26.750 ± 6.956	29.188 ± 5.930
Retweet	#Nodes	87.063 ± 30.005	355.500 ± 107.412	476.813 ± 117.641	720.875 ± 210.047	734.375 ± 153.998
	#Edges	116.375 ± 46.124	722.125 ± 258.400	940.938 ± 277.384	1676.250 ± 693.462	1431.625 ± 351.722
	In/Out degree	0.981 ± 0.102	1.607 ± 0.129	1.653 ± 0.114	1.760 ± 0.143	1.728 ± 0.094
	Density	0.121 ± 0.048	0.009 ± 0.001	0.006 ± 0.001	0.004 ± 0.000	0.003 ± 0.000
	Clustering Coefficient	0.051 ± 0.016	0.078 ± 0.010	0.063 ± 0.008	0.048 ± 0.008	0.060 ± 0.006
	Reciprocity	0.053 ± 0.018	0.066 ± 0.010	0.054 ± 0.005	0.058 ± 0.008	0.070 ± 0.006
	#WCCs	6.250 ± 1.627	6.500 ± 1.780	5.375 ± 1.341	6.625 ± 1.326	6.625 ± 1.998

Table 2: Descriptive statistics (mean ± SE) of network metrics for the retweet and conversation networks over time. Each metric is an average over the individual conferences.

- Combined features: A set of features that utilizes all the features above to test the combination of them.

Evaluation. We used Information Gain to determine the importance of individual features in WEKA [6]. Then we computed the normalized score of each variable’s InfoGain value as its relative importance. To evaluate the binary classifier, we deployed different supervised learning algorithms and used the area under ROC curve (AUC) to determine the performance of our feature sets. The evaluation was performed using 10-fold cross validation in the WEKA machine learning suite.

5. RESULTS

In the following sections we report on the results obtained for each of our analyses.

5.1 Has usage changed?

We answer this question by the results presented in Figure 1, which shows the overtime ratio values of different Twitter usage in the conferences accompanied by their corresponding values from the random dataset.

We can highlight two distinct patterns first. The trends we observed for the information usage ratios are similar. The Retweet Ratio increases (6.7% in 2009, 48.2% in 2013) over the years (one-way ANOVA, $p < .001$) along with URL Ratio (21.2% in 2009, 51.3% in 2013; one-way ANOVA, $p < .001$). Noticeably, Retweet Ratio rapidly increased from 2009 to 2010 but rather steadily gained afterward. We believe this could be explained by the Twitter interface being changed in 2009, when they officially moved ‘Retweet’ button above the Twitter stream [3]. On the other hand, rather stable patterns can be observed in both conversational measures: Reply Ratio (8.2% in 2009, 6.1% in 2013) and Mention Ratio (10.0% in 2009, 12.9% in 2013). Therefore, as we expected, *Twitter behaved more as an information sharing platform during the conference, while the conversational usage did not seem to change over the years.*

Figure 1 also presents the differences between the ratio values in the conference dataset and the baseline dataset, as we want to understand if the trend observed above is the reflection of Twitter usage in general. During the conferences, we observed a higher Retweet Ratio and URL Ratio. We argue that it is rather expected because of the nature of academic conferences: sharing knowledge and their valuable recent work. The Mention Ratio in the conference is slightly higher than it is in the random dataset, because the conference is rather an event of people interacting with each other in a short time period. However, we observe a significant difference in Reply Ratio. Users start the conversation on Twitter using the conference hashtag to some extent like all users, but most users

who reply (more than 90%) usually drop the hashtag. Although it is still a public discussion on Twitter, we guess that users tend to keep it semi-public to their listeners. However, a deeper analysis which is outside the context of this research should be conducted to assess this assumption, since in some cases users would drop the hashtag simply to have more characters available in the message.

5.2 Has interaction changed?

Table 2 shows the evolution of the network measures. Each metric is an average over all the individual conferences in each year from 2009 to 2013. We first highlight the similar patterns over years observed from both networks: conversation network and retweet network. During the evolution, several types of network measures increase in both networks: (i) the average number of nodes; (ii) the average number of edges; and (iii) the average in/out degree. This suggests that more people are participating in the communication network.

Then, we compare these two networks in terms of the differences observed. Table 2 shows that the average number of weakly connected components in conversation network (#WCC) grows steadily over time from 4.750 components on average in 2009 to a significantly larger 29.188 components in 2013, with the CHI conference being the most fragmented (#WCC=87, #Nodes=2117). However, the counterpart value in retweet network is almost invariable, staying between 5.375 and 6.625 on average. The former metric supports the assumption of Ross et al. [19] in terms of the scattered characteristic of the activity network (i.e. multiple non-connected conversations). The #WCC suggests that the retweet network is more connected than the conversation network.

Not surprisingly, the reciprocity in the conversation network is significantly higher than the one in the retweet network ($p < .001$ in all years; pair-wise t-test). This shows that the conversations are more two-way rounded interactions between pairs of users while people who get retweeted do not necessarily retweet back. Both results are rather expected. The mentions and replies are tweets sent to particular users, and therefore the addressed users are more likely to reply back due to social norms. Yet, the retweet network is more like a star network, and users in the center do not necessarily retweet back.

Moreover, we observe that the average clustering coefficient in conversation network is higher than the one in retweet network, in general. We think that two users who talked to the same people on Twitter during conferences are more likely to be talking to each other, while users who retweeted the same person do not necessarily retweet each other. However, the significant difference is only found in 2012 ($p < .05$; t-test) and 2013 ($p < .001$; t-test). We tend to believe that it is the nature of the communication on Twitter,

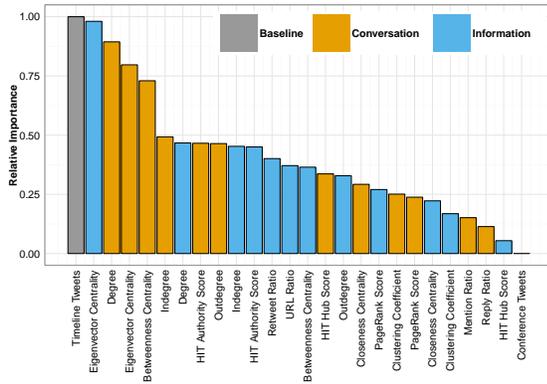


Figure 2: Relative Importance of the features.

but we need more observations and further analysis to support this guess.

5.3 What keeps the users returning?

The results of the prediction model are shown in Table 3, the Bagging approach achieved the highest performance across all the feature sets. Furthermore, baseline features achieved better performance when they are accompanied with conversation features than with the information features. This suggests that the conversation features have more predictive powers in inferring the users’ continuous participation. Finally, the combination of all the features reached the best performance.

We further examine the importance of features in the combined set based on their information gain measures. Figure 2 shows the relative importance of different features. First, it is interesting that two features in the baseline have distinct importance: the number of one user’s conference tweets produced the least information gain, while the number of timeline tweets has the least information gain. This suggests that the user is more likely to return because she has been highly active on Twitter in general; the number of conference tweets tends to be an incident of many factors, for instance, users could have fewer conference tweets because she does not have time to tweet during the conference or the Internet connection is not available.

We observe that the user’s eigenvector centrality score from the information network has a higher importance (second lead in all the features). As the user with a higher eigenvector centrality score means that they and their neighbors have higher values in this measure, we interpret this as these users play important roles in information spreading during conferences. We conjecture these users are influential members in the conference so they are more likely to participate in the future.

Furthermore, we observe that three measures of the conversation network are also prominent: degree centrality, eigenvector centrality, and betweenness centrality. We know from the definitions of these measures that: higher value in the degree centrality of the user means more people this user has talked to in the conference on Twitter; higher eigenvector centrality value suggests the user has conversed with users that are also active in this network; and higher betweenness centrality indicates the user has served as the bridge of different groups of people. We suspect users’ conversational experience and their roles are linked to their decisions of the return in the next year.

6. DISCUSSION AND CONCLUSIONS

In this paper, we investigated how Twitter has been used in academic conferences and how this has changed over the years. We attempted to study the usage, the network structures, and participa-

Feature Sets	RF	ADA	Bagging	LR	SVM
Baseline	0.654	0.713	0.716	0.714	0.63
Baseline+Conversation	0.728	0.754	0.774	0.742	0.659
Baseline+Information	0.74	0.761	0.771	0.74	0.655
Combined	0.753	0.775	0.787	0.754	0.662

Table 3: Area under the ROC curve (AUC) for predicting continuing participation in the coming year with different feature sets and learning algorithms. The best algorithm for each feature set is highlighted. Methods used are Random forest (RF), Adaptive Boosting (ADA), Bagging, logistic regression (LR), and support vector machines (SVM).

tion of the users during conferences on Twitter. We addressed our research with three questions.

To answer the first research question *RQ1*, we computed four features to quantitatively measure two aspects of Twitter use at conferences: information usage and conversation usage. Our results show that researchers are using Twitter in 2013 in a different way than they did in 2009 with respect to favoring more the information sharing value of Twitter, however their Twitter conversation interactions in 2013 do not differ much from five years ago.

Then, to answer the second research question *RQ2*, we constructed the conversation network and the retweet network for each conference and use network properties to measure how people interacted over the years. Our results show that with more people participating over time, the conversations turn scattered into small groups, while the information flow (retweets) stays mainly within a giant component.

Finally, to answer the third research question *RQ3*, we trained a binary classifier with features extracted from the Twitter usage, network positions of the user and achieve the best prediction performance using the combination of conversation features, information features, and baseline features. We also found that the most influential factors that drive users’ continuous participation are being active on Twitter, being central in the information network and talking with more people.

We acknowledge the limitations of this work. First of all, we claimed to obtain the quasi-random tweets from the historical tweets, however this approach still needs refinement on the choices of several parameters, although we achieved reasonable results compared to other studies³. Second, we assumed conference hashtags as the medium with which users shared information and interacted with others during and about the conference. Although this is adopted by related studies, we may have collected incomplete data because some tweets that do not contain the conference hashtag however still are about the conference. Last but not least, we conducted our analysis on 16 Computer Science conferences, but the findings might not be generalized to all the computer science conferences since users in different sub-domains may show different Twitter usage. In future work, we plan to extend our study to a larger set of conferences across different domains in the field such as Information Systems and Data Mining, in order to see whether users in these conferences behave differently in Twitter.

Acknowledgments: This work is supported by the Know-Center and the EU funded project Learning Layers (7th Framework Program, Grant Agreement 318209). The Know-Center is funded within the Austrian COMET Program - Competence Centers for Excellent Technologies - under the auspices of the Austrian Ministry of Transport, Innovation and Technology, the Austrian Ministry of Economics and Labor and by the State of Styria. COMET is managed by the Austrian Research Promotion Agency (FFG).

³As of the time we wrote this paper, we did not find other statistics for random sample that can serve as a comparison with our study. However, later, Liu et al. [14] provided statistics in their study and their results similar to ours.

7. REFERENCES

- [1] F. Abel, C. Hauff, G.-J. Houben, R. Stronkman, and K. Tao. Twitcident: fighting fire with information from social web streams. In *Proceedings of the 21st international conference companion on World Wide Web*, pages 305–308. ACM, 2012.
- [2] P. Bonacich. Factoring and weighting approaches to status scores and clique identification. *Journal of Mathematical Sociology*, 2(1):113–120, 1972.
- [3] D. Boyd, S. Golder, and G. Lotan. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *System Sciences (HICSS), 2010 43rd Hawaii International Conference on*, pages 1–10. IEEE, 2010.
- [4] T. Desai, A. Shariff, A. Shariff, M. Kats, X. Fang, C. Christiano, and M. Ferris. Tweeting the meeting: an in-depth analysis of twitter activity at kidney week 2011. *PLoS one*, 7(7):e40253, 2012.
- [5] M. Ebner. Introducing live microblogging: How single presentations can be enhanced by the mass. *Journal of research in innovative teaching*, 2(1), 2009.
- [6] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: An update. *SIGKDD Explor. Newsl.*, 11(1):10–18, Nov. 2009.
- [7] H. He and E. A. Garcia. Learning from imbalanced data. *IEEE Trans. on Knowl. and Data Eng.*, 21(9):1263–1284, Sept. 2009.
- [8] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *JOURNAL OF THE ACM*, 46(5):604–632, 1999.
- [9] J. Lanagan and A. F. Smeaton. Using twitter to detect and tag important events in live sports. *Artificial Intelligence*, pages 542–545, 2011.
- [10] A. O. Larsson and H. Moe. Studying political microblogging: Twitter users in the 2010 swedish election campaign. *New Media & Society*, 14(5):729–747, 2012.
- [11] J. Letierce, A. Passant, J. Breslin, and S. Decker. Understanding how twitter is used to spread scientific messages. 2010.
- [12] J. Letierce, A. Passant, J. G. Breslin, and S. Decker. Using twitter during an academic conference: The# iswc2009 use-case. In *ICWSM*, 2010.
- [13] Y.-R. Lin, B. Keegan, D. Margolin, and D. Lazer. Rising tides or rising stars?: Dynamics of shared attention on twitter during media events. *CoRR*, abs/1307.2785, 2013.
- [14] Y. Liu, C. Kliman-Silver, and A. Mislove. The tweets they are a-changin’: Evolution of twitter users and behavior. 2014.
- [15] M. Mendoza, B. Poblete, and C. Castillo. Twitter under crisis: can we trust what we rt? In *Proceedings of the first workshop on social media analytics*, pages 71–79. ACM, 2010.
- [16] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web, 1999.
- [17] W. Reinhardt, M. Ebner, G. Beham, and C. Costa. How people are using twitter during conferences. *Creativity and Innovation Competencies on the Web*, page 145, 2009.
- [18] B. Rieder. The refraction chamber: Twitter as sphere and network. *First Monday*, 17(11), 2012.
- [19] C. Ross, M. Terras, C. Warwick, and A. Welsh. Enabled backchannel: conference twitter use by digital humanists. *Journal of Documentation*, 67(2):214–237, 2011.
- [20] A. Sopan, P. Rey, B. Butler, and B. Shneiderman. Monitoring academic conferences: Real-time visualization and retrospective analysis of backchannel conversations. In *Social Informatics (SocialInformatics), 2012 International Conference on*, pages 62–69. IEEE, 2012.
- [21] A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welpe. Predicting elections with twitter: What 140 characters reveal about political sentiment. *ICWSM*, 10:178–185, 2010.
- [22] S. Wasserman and K. Faust. *Social network analysis: methods and applications*. Cambridge University Press, Cambridge; New York, 1994.
- [23] K. Weller, E. Dröge, and C. Puschmann. Citation analysis in twitter: Approaches for defining and measuring information flows within tweets during scientific conferences. In *Proceedings of Making Sense of Microposts Workshop (#MSM2011). Co-located with Extended Semantic Web Conference, Crete, Greece, 2011*.
- [24] X. Wen, D. Parra, and C. Trattner. How groups of people interact with each other on twitter during academic conferences. In *Proceedings of the Companion Publication of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW Companion '14*, pages 253–256, New York, NY, USA, 2014. ACM.
- [25] A. Zubiaga, D. Spina, V. Fresno, and R. Martínez. Classifying trending topics: a typology of conversation triggers on twitter. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, pages 2461–2464. ACM, 2011.