TwitterSigns: Microblogging on the walls

Markus Buzeck
IfGI, University of Münster
Weseler Str. 253, 48151 Münster
Germany
markus.buzeck@uni-muenster.de

Jörg Müller
Deutsche Telekom Laboratories, TU Berlin
Ernst-Reuter-Platz 7, 10587 Berlin
Germany
joerg.mueller@tu-berlin.de

ABSTRACT

In this paper we present TwitterSigns, an approach to display microblogs on public displays. Two different kinds of microblog entries (tweets) are selected for display: Tweets that were posted in the immediate environment of the display, and tweets that were posted by people associated with the location where the displays are installed (locals). The prototype was tested in a university setting on 4 displays for 4 weeks and compared to the information system that is usually running on the displays (iDisplays). Using face detection we show that people look significantly longer at TwitterSigns than at iDisplays. Interviews show that the relationship of viewer and poster as well as the tweet content are much more important than time and location of the tweet. Viewers recall and recognize mostly tweets from people they know, and of apparent importance for themselves (like a apparent bomb found in the city center). Furthermore, TwitterSigns change the way people use twitter (e.g. they feel more responsible for what they tweet). Passers-by seem only to look for keywords and only stop and read the whole tweet if they found some interesting keyword.

Categories and Subject Descriptors

H.5.1 [**Multimedia Information Systems**]: INFORMATION INTERFACES AND PRESENTATION

General Terms

Design, Human Factors

Keywords

Public Displays, Microblogging, Field Study

1. INTRODUCTION

The development of new display technologies, like e-paper and OLED, may soon lead to 'electronic wallpaper' covering many surfaces in buildings and urban areas. Such displays have the benefit that they are naturally local, thus information about the immediate environment may suit them well. In contrast to static signs, they can also be updated often and can thus show information about what is happening 'right here, right now'. So they could improve people's situation awareness by giving them access to more information about their immediate surroundings than are available to their normal senses. One possibility to obtain this information about the immediate environment are microblogs. The case of an airplane crashing into Hudson river for example was first reported on Twitter[1], and showing such information on public displays in the surroundings can make people immediately aware of such events.

Evaluating whether such applications for public display actually work as expected however is a difficult problem. The main difficulty is that interactions of viewers with the public displays are usually very *sparse* and *short*. For one display, there may be only about 10 viewers a day, each looking only for about 1-2 seconds. This makes formal observations of interactions as well as laboratory studies practically impossible. Additionally, viewers can remember almost none of their interactions with the displays, making also interview techniques very difficult.

In this paper, we use a deployment of a public display system in an area with many natural passers-by, thereby generating a big number of very short interactions with different people in their naturally occuring social context. Second, we use automated audience measurement via a face detection system to measure when the audiences faces are turned towards the displays. We select a sample of the audience population to ask them what content from the displays they recall and present them a sample of content items to test their recognition. We augment this quantitative data with qualitative semi-structured interviews and casual observations, as well as analysis of other data sources (their Twitter behavior).

2. RELATED WORK

The first system to show Twitter information on public displays is Twitterspace by Hazlewood et al. [1]. Twitterspace focuses exclusively on awareness within a local community and shows the 100 most recents tweets of Twitter users following a central community account. It provides a sense of *community-at-a-glance* and tweets are sorted by posting time on a single big display in a community space. Java et al. [2] have identified intentions to use Twitter such as *daily chatting* (talking about daily routines and current

[1]http://twitpic.com/135xa

activities), *conversation* (comment or reply friends news using the @ character), *sharing information* (containing shortened URLs), and *reporting news* (like latest news or weather through automated agents or RSS feeds). Building on this work, Krishnamurthy et al. [3] have categorized Twitter users by their follower/following behaviour. *Broadcasters* like radio stations follow less people but have many followers and publish information, e.g the current played song. *Acquaintances* have almost the same number of followers as they themselves are following. The last group consists of *miscreants* that try to follow as many people as they are allowed to. Miscreants have only very few followers and often only one published tweet with a redirecting URL.
GroupCast [5] is one example of public displays intended to enhance community awareness, and ReflectiveSigns [6] is one example how cameras and face detection can be used to measure the audience of public displays.

3. TWITTERSIGNS

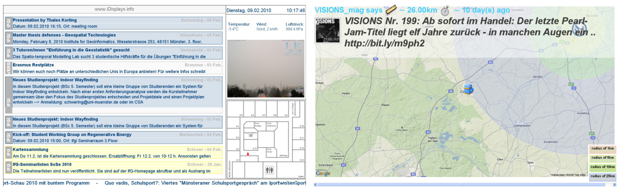

Figure 1: Left: The information system usually running on the display network (iDisplays [7]) . On the left part, information from faculty is shown, while on the right part, data about the environment (e.g. weather information) are presented. Right: Screenshot of TwitterSign. One tweet at a time is presented in front of map that shows the location from where it originated and a photo of the author.

The intention of TwitterSigns was to increase situation awareness by showing what is happening 'right here, right now' on public displays. The microblogging platform Twitter serves as a content source, and content is retrieved from Twitter in two different ways. First, a dedicated Twitter account is created for the displays that follows all institute employees and students who are known to use Twitter. Second, Twitter is searched for all recent tweets that have been posted in a 25 km radius around the display. Because the displays are installed in transitional spaces, where most people only pass by, we expected that most passers-by would only look at the displays for about 2 seconds. Therefore, at a given time, only a single tweet is presented for a time slot of 10 seconds (see figure 1). In order to emphasize the location where the tweet was posted, a corresponding map is shown along with the tweet, as well as how long ago the tweed was posted and a photo of the author. At any given moment, a random tweet is presented, with the probability of a tweet being shown determined by whether it is posted by a local, how old it is, and how far away from the display it was posted. The weight for tweet t is currently calculated by $w(t) = f_1(t)\frac{1}{f_2(t)}\frac{1}{f_3(t)}$, where $f_1(t)$ is 5, if the tweet is from a local, 3 if it is from the university, and 0.1 otherwise, $f_2(t)$ is the age and $f_3(t)$ the distance between the display and the tweet, respectively. All of these parameters were tweaked iteratively based on our own impression and feedback from users.
TwitterSigns is implemented as a server and a display component. The server component is a PHP script which connects to Twitter using Twitters REST and SEARCH APIs and retrieves the 20 most recent tweets from a 25km radius around the display, as well as all tweets posted by locals. It also conducts some filtering for assumingly uninteresting tweets (which start with RT or the @ symbol). The display component is a combined PHP / JavaScript script which then calculates the weight for each tweet and randomly selects one of them for display. The tweet position is geolocated and Google Maps is used to show the location of the tweet along with the tweet itself.

4. METHOD

We deployed TwitterSigns on four public displays in a university department for four weeks from September 14, 2009 to October 11, 2009. The displays are usually running the iDisplays university information system which is in use since October 2005 [7] . Display 1 is installed in the entrance of the institute, display 2 in a sofa corner, display 3 in a coffee kitchen, and display 4 in a corner of a hallway. The audience consists of institute employees, students and visitors. In order to compare TwitterSigns to the university information system, content was rotated. In weeks 1 and 2, TwitterSigns was shown on displays 2 and 4 and the information system was shown on displays 1 and 3, whereas in weeks 3 and 4, TwitterSigns was shown on displays 1 and 3, whereas the information system was shown on displays 2 and 4. During the four weeks of deployment, 21,099 different tweets were shown on TwitterSigns.
We equipped the displays with cameras and use face detection [4] to determine if, and for how long, there is a face turned towards the display. For one day we recorded videos of people passing the displays, correlated with the face detection, to ensure there is a relatively high correlation between faces seen by the camera and people actually looking at the screen. Timestamp, location, as well as duration for which the face was detected, are all stored to a central database.
In addition to the face detection data, we conducted recall/recognition tests and semi-structured interviews with 8 students and institute employees in the days after the four week deployment period. Participants were selected on an opportunity basis. For the recall tests we simply asked the participants for all information shown on the displays which they could remember. For the recognition tests we printed out TwitterSign screenshots of 6 tweets that actually appeared on the displays and asked participants whether they could remember these. In the semi-structured interviews we asked for peoples usage and opinion of TwitterSigns. In addition, we conducted informal observations of behaviour on the hallways during the deployment period.

5. RESULTS

During the deployment period, we detected 9,575 individual views towards the displays. We present results of the analysis of view times, as well as from the interviews.

People look significantly longer at TwitterSigns.

The first step of data analysis was to look at the individual view times, that is, if a face was found in front of a display, how long it was there. People look significantly longer at

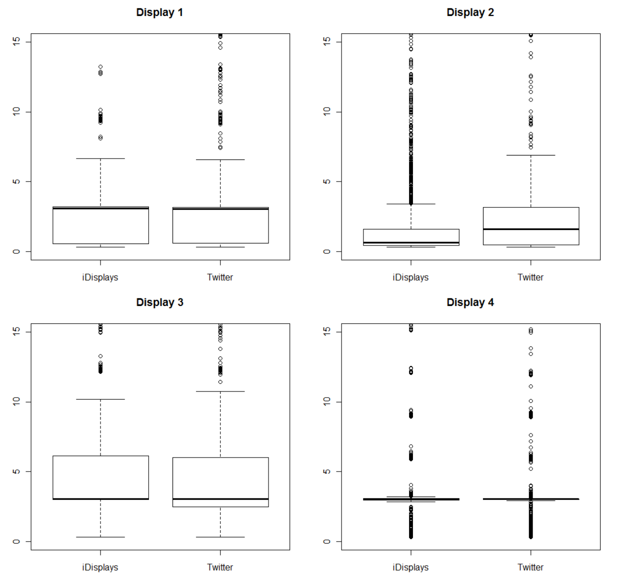

Figure 2: Boxplots of view times towards the four displays. For all displays except display 3, people looked significantly longer when TwitterSign was displayed.

TwitterSigns than at the university information system for all displays except the one installed in the kitchen. The view times for the individual displays are shown in figure 2. We see that for the display in the entrance, average view time was 2.96 sec. for the iDisplays case, and 3.53 sec. for the TwitterSigns case. This difference is significant[2]. For the display in the sofa corner, average view time for iDisplays was 2.03 sec., and 3.41 sec. for TwitterSigns. Again, this difference is significant[2]. For the display in the hallway, average view time for iDisplays was 3.35 sec., and 3.77 sec. for TwitterSigns (significant[2]). For the display in the coffee corner, average view time was 5.09 sec. for iDisplays, and 5.57 sec. for TwitterSigns. This difference however was not significant.

While TwitterSigns seem to attract more attention overall, this is not significant.

The second step of analysis was to look not only at the individual view times, but at the overall attention that was attracted by the different content. Therefore, for each display and each day of the study, the total view time for that day was calculated. For the display in the entrance, average view time per day for the iDisplays case was 235.22 sec., while it was 359.49 sec. for the TwitterSigns case. For the display in the sofa corner, average view time per day was 483.91 sec., while it was 181.45 sec. for the TwitterSigns case. For the display in the hallway, average view time per day for iDisplays was 352.33 sec., while it was 493.33 sec. for TwitterSigns. For the display in the kitchen, average view time per day for iDisplays was 234.48 sec., while it was 285.29 sec. for TwitterSigns. None of these differences were significant.

[2]ANOVA, $p < .05$

People recall and recognize tweets from people they know, about an assumed bomb on the town square, and about the university.

While the face detection data tells us that people have looked at the displays, it does not tell us whether they remember any of the content. The recall interviews showed that compared to the view times at the displays, recall was surprisingly low. Participants could only recall one or two tweets (on average 1.4), and even of these mostly only the vague content, not the exact message. Looking more closely at what kind of tweets were recalled shows some interesting pattern: P1 recalled to have seen press annoucements of the university, P2 weather forecast, P3 recalls two tweets from locals, P4 doesn't recall any tweets, P5 and P6 recall a tweet about an assumed bomb on the town square, P6 also recalls having seen his own tweets on the displays, P7 recalls one tweet from university sports and one from somebody having a new tattoo, and P8 recalls one tweet of a university employee and one of a scientist who made some discovery.
It is striking that compared to recall, recognition was relatively high. One tweet about a university sports hall being dismantled was recognized by 4 participants, two tweets by university employees (one working on somebodies diploma thesis, one going home after work) were recognized by 3 participants each, and three tweets from random people living nearby (back from sports, a sparking notebook, and a conversation-tweet about offensive tweets) were not recognized by any participants.

Tweets are interesting when you know the poster or the content is interesting, not when they are new or close.

All interviewees stated that they considered most of the tweets of people from the surroundings that they had no relation to as trash and were not interested in them. "general microblogging problem. Much of what is posted there is just blabla, trash or things which I am not interested in" (P3). They stated that mostly tweets from people they know or which have some content which affects them personally are interesting, and time and location of the tweet play a much minor role in comparison "If the information is not especially related to the location, I consider the location quite irrelevant" (P8). In particular, many participants emphasized that the location of a tweet only plays a role only in very specific circumstances, for example for the weather forecasts or to see where people they know currently are. Location in general was only considered interesting if it is an exact coordinate and not only the general area of a city "somewhere in iDisplays, and the coordinate doesn't exist, it is less interesting than if it has some real coordinate" (P1). Similarly, participants stated that time only plays a role for certain tweets. For example, some tweets like weather forecast are quickly outdated and should therefore automatically be removed "If it was up-to-date one time and isn't anymore, it should be deleted" (P1). Similarly, time plays a role for event annoucements.

People adapt to the additional exposure by changing their Twitter behavior.

During the deployment we could observe some very interesting changes in Twitter behaviour, on which the interviews revealed more details. One apparent change was that one local very active on Twitter (P6) suddenly changed his Twitter image to a photo of somebody else. The interview revealed that one evening he had written a tweet about going to bed after playing a game on his iPhone until his thumb hurts. Coming to work next day he saw his tweet up on the displays and realized that it could be seen by the students and leave a somewhat negative impression on him. In order not to be recognized by people who don't know his Twitter name, he changed his photo to the photo of somebody else. Two participants spontaneously started using Twitter during the TwitterSigns deployment. While one of them stopped after posting only 3 tweets, the other used it continuously during the deployment posting 21 tweets.

People seem to scan for keywords while passing by, stopping only if they see an interesting keyword.

Looking at the view times, it becomes apparent that most passers-by tend to look at the displays for less than four seconds. This is supported by our casual observations that most people look very shortly at the displays while passing by without stopping. Only very rarely somebody stops in front of a display to read the content. Even for the very short Twitter messages, it seems unlikely that passers-by can read the whole tweet in this short time span. This was supported by the interviews. P5 for example stated that he only scans the tweets for keywords, and only stops to read the whole tweet if some keyword catches his interest: "You just scan over it, and if you find some interesting words, then you look more closely" (P5). Another participant stated that he first looks at the photo of the author, and does't read the tweet at all if the photo seems dubious.

6. DISCUSSION

While the face detection showed that in general, the TwitterSigns make people look longer than the university information system, much more detail was delivered by the other methods. In particular only tweets from people the participants knew were recalled and recognized, as well as tweets that affected them personally. In the interviews, the participants stated that they considered many of the other tweets as trash. Using just one of the methods alone, one might have had a very different impression. Either that TwitterSigns are successful, independent of the content (using only face detection), or that nobody would look at them at all (using only the semi-structured interviews). The basic assumption that we started out with was that it could be interesting for people to know what is happening 'right here, right now'. This seemed not to be the case. The participants considered most of the tweets presented as trash, and time and location do not seem to be a good filter to find interesting tweets. Much more work needs to be done to filter out the interesting tweets, with relationship to the viewer, as well as organizational relationship to the organization where the display is installed, and inherent news value being the most important directions. Such tweets were indeed considered interesting by the participants. One very important experience from this study was that view times need to be normalized by footfall. We saw that the average view time for TwitterSigns was significantly longer than for iDisplays, while this was not the case for the total view time per day. Looking closer at the average view times per day revealed why. There we see an enormous spike of view time for the displays in the sofa corner and the entrance for Wednesday, Oct. 7. Interviews with institute employees revealed that there was a tour of the institute for freshmen on that Wednesday, and that they had to wait around the sofa corner and the entrance for quite some time. This shows that there can be enormous differences in the number of passers-by (footfall), which is caused by external events, and completely independent from the content. In order to measure the impact of the content on attention, it would be advisable to measure the footfall independent from attention (as is partially achieved by looking at the average view times).

7. CONCLUSION

In this paper we have presented TwitterSigns, public displays that show microblog entries from the immediate environment in order to improve situation awareness. Passers-by look significantly longer at the TwitterSigns than at the university information system that usually runs on the displays. However, they remember only very specific tweets from people they know or which could affect them personally, and are also not interested in the others. The next versions of TwitterSigns should therefore concentrate at identifying which tweets are interesting to the current viewer, for example by identifying his friends or evaluating the news value of tweets in general. In this study it proved useful to combine quantative data from face detection with qualitative interview data. Using face detection however, footfall should be measured independently from attention, to account for strong external effects that influence attention independent from the application.

8. REFERENCES

[1] W. R. Hazlewood, K. Makice, and W. Ryan. Twitterspace: A co-developed display using twitter to enhance community awareness. In *PDC '08: Proceedings of the Participatory Design Conference*, 2008.

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

[3] B. Krishnamurthy, P. Gill, and M. Arlitt. A few chirps about twitter. In *WOSP '08: Proceedings of the first workshop on Online social networks*, pages 19–24. ACM, 2008.

[4] C. Küblbeck and A. Ernst. Face detection and tracking in video sequences using the modified census transformation. *Image and Vision Computing*, 24 (6):564–572, 2006.

[5] J. F. McCarthy, T. J. Costa, and E. S. Liongosari. Unicast, outcast & groupcast: Three steps toward ubiquitous, peripheral displays. In *Ubicomp '01: Proceedings of the 201 Conference on Ubiquitous Computing*, pages 332–345, 2001.

[6] J. Müller, J. Exeler, M. Buzeck, and A. Krüger. Reflectivesigns: Digital signs that adapt towards audience attention. In *Proceedings of Pervasive 2009, Nara, Japan 2009.*, volume 5538 of *Lecture Notes in Computer Science*, pages 17–24. Springer, 2009.

[7] J. Müller, O. Paczkowski, and A. Krüger. Situated public news and reminder displays. In *AMI '07: Proceedings of the 2007 European Conference on Ambient Intelligence*, pages 248–265, 2007.