SPOTSENSING

A Plataform for Social Web Sensing

Nuno Correia, Filipe Meneses and Rui José
Centro Algoritmi, Universidade do Minho, Guimarães, Portugal.
pg13107@alunos.uminho.pt, meneses@dsi.uminho.pt, rui@dsi.uminho.pt

Abstract—**Social Networking Sites (SNS) have an unprecedented ability to capture Human activity, including information about the specific physical settings in which those activities are taking place. This represents a major potential for uncovering, on a large scale, new knowledge about aggregate behaviors in the use of places. In this paper, we explore the concept of social web sensor, as a systematic data collection process that can be virtually attached to a particular location to retrieve location-based information from social network sites. This process is completely based on geographically scoped queries to SNS APIs and does not depend on real physical sensors. The objective of this study is mainly to assess the viability of this concept and uncover the potential and limitations of this approach as a reality mining tool for urban environments. We have created an initial implementation and conducted the respective evaluation through the deployment of a number of sensors in the city of London and the analysis of the respective results.**

Keywords-component; Social Sensors, Web 2.0, Social Networks, Twitter, Web Activity.

I. INTRODUCTION

The emergence of ubiquitous computing and Social Networking Sites (SNS) means that digital presence is increasingly shaping our lives and our actions. These services have been capturing Human activity in unprecedented ways, and, in that process, they are leading to the extensive generation of new content associated with those activities. The information includes not only the primary social object, but also meta-data about the circumstances in which that content is generated.

Interestingly, many of these actions are increasingly traceable to the specific physical settings in which they occur. For example, many twitter posts are now tagged with location information and emerging systems, like Foursquare or Gowalla, are all about signaling presence in places. If we also consider that much of this information is publicly available, we may easily acknowledge that there is a clear potential for using these tools for uncovering, on a large scale, new knowledge about aggregate behaviors in the use of places.

In this paper, we explore the concept of social web sensor, as a systematic data collection process that can be virtually attached to a particular location to retrieve location-based information from social network sites. Through geographically oriented queries, a social web sensor obtains data generated at a particular location and makes it available, possibly in some aggregate form that constitutes the sensor output. This process is completely based on user-generated data and without directly recurring to any physical sensor. Still, the increasing use of geo-tagged or geo-referenced data may provide the ground for making these sensors valuable tools for detecting and understanding location-based events and patterns, particularly in urban environments. Our objective is thus to study the viability of this concept and uncover the potential and limitations of this approach as a reality mining tool for urban environments.

In Section II, we revise related work and clarify the relationship between our contribution and previous work on this topic. In Section III, we present the Spotsensing architecture for creating location-based social sensors. To evaluate the concept, we have conducted a sensing study using a number of sensors deployed in the city of London. The description of this evaluation and the analysis of the results obtained is included in Section IV. Finally, in Section V, we present our concluding remarks and outline future work.

II. RELATED WORK ON SENSING AND SOCIAL ANALYSIS

The ubiquity of mobile phones, GPS-enabled devices, CCTV, Wi-Fi networks and other sensors provides many opportunities for collecting data about Human activities. The collection of that information on a large scale provides the opportunity to obtain an increased understanding of the dynamics of communities and uncover new knowledge about the rhythms of city life and the digital fingerprint of the urban environment. The collaborative use of physical sensors has been explored in several projects as an approach for inferring patterns of Human behavior, such as CenceMe [1] and MetroSense [2]. CenceMe uses audio sensors and accelerometers to be able to "read" the environment around the user and MetroSense explores the interaction between the sensors carried by the user and other devices available (e.g. Wi-Fi access points, etc.). Real Time Rome [3], an MIT project developed in 2006, aggregates data from mobile phones, buses and taxis in Rome to better understand urban dynamics in real time. Collected data was used to create a set of visualization plots exposed in the 2006 Venice Biennale. SensorPlanet [4] is a global research framework for mobile device centric wireless sensor networks. It enables collaborative sensing initiatives on a large scale and includes a central repository for sharing the data. The Intel-sponsored Urban Atmospheres project [5] aims to understand our future evolution of digital and wireless computing and how this will influence, disrupt, expand and integrate into social patterns existent within our public urban landscape. The MIT Cartel project [6] was designed to collect, process, deliver and visualize urban-sensing data about traffic conditions. While these approaches may have an important role in our ability to sense physical environments, they require the physical presence of the sensors and they tend to generate lowlevel data, typically associated with movement, sound, air quality or temperature.

The use of data obtained from social networks to understand Human activities in physical space has also been explored in a number of projects. Data from personal location disclosures, even when aggregate, may provide a level of social awareness that can be used for that purpose. Foursquare and Gowalla are both location-based social networks that encourage people to explore cities. In Foursquare, points are won by sending status messages with the user's location and in Gowalla users can collect virtual objects left in places they visited [7]. All these system generate as part of their usage, considerable data about usage patterns of physical spaces. HomeAndAbroad is an example of how to leverage on photos published in Flickr for trip planning [8]. By using the location and time of the photos taken on the city, the system automatically generates gazetteers that include the average time to spend at the main attractions in four main cities. Several other projects have explored the use of Twitter data to infer new data about the cities and about the users. Tweet-o-Meter [9] shows, almost in real time, the number of tweets being published in different cities. Although the population percentage that uses Twitter in each city may be different, the visual rendering of tweet-o-meter provides a descriptive view of which cities are more active and detects when the activity grows or decreases. Fujisaka et al. describe a method for detection of unusual crowded locations by monitoring micro-blogging systems to obtain geo-tagged data which is analyzed to examine spatial distribution [10]. TweetHood [11] predicts the location of the user based on location of his closest friends. Although several projects have used the data from social networks to uncover information about the physical world, there is no simple and straightforward solution to generate such data for a particular physical location. In this work, we focus specifically on how to create generic and multi-purpose data collection processes that may represent the most significant elements of the social activity associated with a location.

III. THE SPOTSENSING ARCHITECTURE

This section describes the SpotSensing architecture and how it can be used to deploy location-based web sensors.

A. Location-based web sensors

A location-based social web sensor is an abstraction for a location-based data collection process that produces a specific output. The sensor metaphor is used to convey the idea that those processes are attached to specific locations, just like any sensor. However, these sensors differ from traditional sensors in two fundamental ways: the first is that they are exclusively virtual, which means that their attachment to a location is not the result of their physical presence at that location, but simply a criterion in the data collection process; the second is that they do not sense physical variables, such as temperature, speed, light, or any other. They are based on a set of data sources corresponding to existing social networking platforms. This is where the data comes from and potentially any type of social network may be used, provided there is some type of support for location-based queries. Sensors will therefore be measuring social activity as represented in observable behavior within social networks. The social networking services from which data is collected are called Data Sources.

A sensor is therefore a virtual entity created by clients of the SpotSensing infrastructure to measure social activity associated with a particular location. A sensor will always have a type, which determines the type of data it generates, and a coverage area, which determines the geographical scope upon which it is expected to generate data.

A sensor type is not determined by the data source from which it obtains the information, but by its output data. A particular sensor type could combine data from multiple data sources, and multiple sensor types could provide different views of a particular data source. However, in this initial study we have created two basic usage sensors for Twitter and Flickr. In this paper, we will only analyze in more detail the operation of the twitter sensor, but having created two different types of sensors has been important in highlighting some of the implications of the respective data models to the conceptualization of the sensors. We will refer some of those issues in the conclusions.

Regarding the geographical scope, sensors can be specified as individual sensors or as sensor grids. The geographical scope of an individual sensor is specified by setting the coordinates of the center and a radius. The geographical scope of a sensor grid is specified by setting the coordinates of the center, a radius for the grid, and the number of sensor levels inside the grid. The system will automatically generate a coordinated and multi-level set of sensors covering with minimal overlap the specified region, as exemplified in Figure 1.

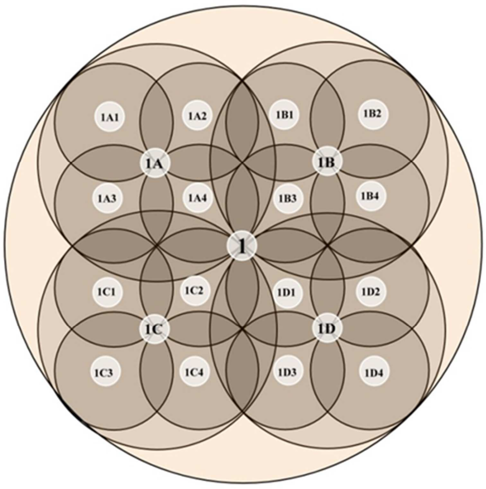

Figure 1 - Sensor Grid

The image represents a larger sensor, called sensor 1, containing four other sensors 1A-1D, with a 50% radius, and then 16 other smaller sensors, e.g. 1A1, with 25% the radius of the larger sensor.

B. SpotSensing Service

The architecture of the SpotSensing service is depicted in Figure 2.

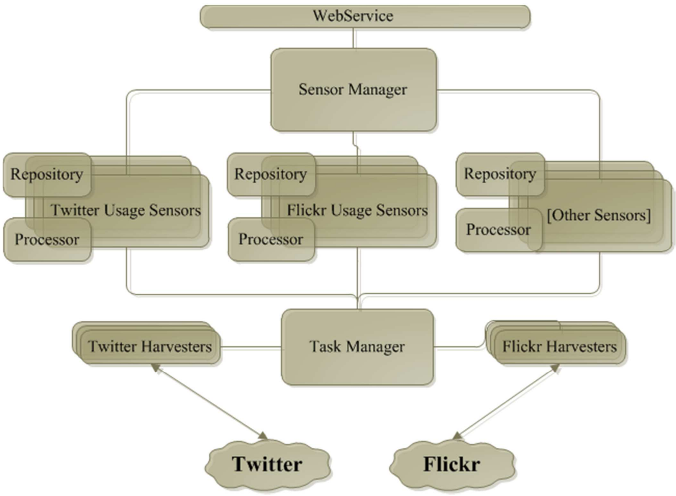

Figure 2 – Architecture of SpotSensing

A web service interface enables clients to request the creation of new sensors and obtain the data they generate. A sensor manager handles all the information about existing sensors and controls their life cycle. The sensing process itself is controlled by sensor modules corresponding to the sensor types that have been requested. A sensor types encapsulates the data processing knowledge associated with a particular form of social sensing. In our initial implementation, we have developed a twitter usage sensor, which will be described in more detail, and a Flickr usage sensor type that obtains data about the images published at a given location and uses the rich meta-data associated with these images, e.g. type of machine used, time elapsed between taking the photo and publishing, tag density, authors diversity and others to generate a number of metrics about photo practices that may also be relevant indicators about the nature of places. Each sensor type knows how to generate the appropriate sensing tasks and how to store the respective data. A Data Repository represents the data collected for a particular type of data source and it is from this data that the sensor output will be generated.

The Task Manager accepts sensing tasks from sensors. A sensing task specifies a periodic data collection process to be conducted on a specific data source. Whenever a new sensor is activated, it needs to generate the sensing tasks that will satisfy its data needs. A Data Harvester is the element responsible for generating the queries to the data source according to the appropriate periodicity. It optimizes the queries associated with a particular data source and periodically retrieves from those data sources the data needed for the currently active sensors. A single sensing task may create multiple Data Harvesters to distribute the load or overcome any limits imposed by data sources.

C. Twitter Usage Sensor

The Twitter Usage Sensor is one of the two types of sensor that we have implemented as part of this study. It supports the generation of information about Human activity in a specific location, as seen from twitter activity.

Twitter may be classified as a micro-blogging tool in which people can share individual messages not exceeding 140 characters. Given the limitations in message size, microblogging is particularly suited for a faster mode of communication in which messages are frequently produced as part of daily life. With the growing popularity of mobile clients for twitter, location information can be easily recorded, and thus an increasing number of tweets are generated with location information as meta-data. The messages may also include a set of codes from which additional information can be inferred, such as the status or mood of the sender or the message target.

The creation of a Twitter usage sensor in the SpotSensing system can be accomplished by specifying the following parameters.

- The coordinates where the sensor is to be deployed.
- The radius of the coverage area of the sensor
- The time granularity to be used for data update.
- A Location Mode that defines whether the sensor should consider only tweets with geographic coordinates or whether it should also consider tweets associated with location through location names.

Access to the twitter API is conducted by specialized twitter harvesters. The Twitter API is based on a RESTful paradigm and provides access to geo-referenced tweets, including all their meta-data, and respective authors.

The Twitter API enforces some limitations on the request rate and number of results that can be obtained. The limitations are based on the IP address of the requester and limit the number of requests to 350 per hour. The API requests may include a definition of the number of tweets to return per page, up to a max of 100, and the page number (starting at 1) to return, up to a max of roughly 1500 results (meaning 15 pages and 100 results per page). Regarding location scoping, the search requests may specify a radius parameter with a minimum of 1 Km.

Multiple harvesters may be used collaboratively in the context of a single sensing task. The periodicity with which they will contact the API is determined by the time granularity specified for the sensor, but also by its coverage radius. An harvester collecting data for a large sensor, will have to make more frequent connections to the data source in order to avoid limitations on the number of tweets retrieved from the API.

The information collected by twitter harvesters is then processed to extract the information that will be exposed by the sensor. The tweets themselves are not part of that information and are only kept for the time needed to process them.

The information exposed by this Twitter usage sensor, includes generic sensor information and usage metrics. The sensor data includes the date of the sensor's deployment, the sensor area in Km^2, the total number of days and the total number of hours in which data was sensed. The usage metrics for the sensing period are as follows:

- The number of tweets;
- The number of tweets sent from one user to another (i.e. dialogs);

- The number of tweets written in English;
- The number of unique users generating tweets;
- The number of followers (i.e. users that add another users as their friends following their tweet's updates);
- The number of tweets sent from specific devices, e.g iPhone or Blackberry;
- The number of tweets with specific emotions tags, links and Retweets (i.e. tweets with the tag "RT", referring someone's tweets).

IV. EVALUATION

The evaluation of the system was conducted through the virtual deployment of a number of twitter-based sensors and the analysis of the data generated by those sensors. The main objectives of this evaluation were as follows:

• To assess the viability of twitter-based sensors in terms of the respective data density and granularity;

• To assess the reliability of the location information associated with tweets;

• To assess the ability of Twitter Usage Sensors to produce distinguishing data across different types of places.

A. Sensor Deployment

A key decision in our study was where to deploy the sensors. Given that different locations may have very different patterns of twitter usage, the chosen location could have a major impact on the conclusions of the study. We have chosen to go for a location that could maximize the number of tweets generated per location, and thus chose a large city, more specifically London. This approach presents two main advantages: The first is that it makes the bias of the location choice more clear, i.e. whatever results we get, they are clearly for highly urban settings with highly connected populations. The second advantage is that this gave us the opportunity to explore the sensors where they are more likely to be effective and therefore generate more relevant findings.

A second decision was how many sensors to deploy, their size and their dispersion. Considering the multiple objective of the study, we have chosen to deploy two sets of Twitter-based sensors. The first set was a sensor grid as represented in Figure 1. This grid includes a total of 21 sensors carefully placed to guarantee the containment of the smaller sensors into the larger ones. We thus had 16 small sensors with a 2 kilometers radius. Each group of 4 of those small sensors was contained within a medium-sized sensor with a 4 kilometers radius. These 4 medium-sized larger sensors were in turn contained within a large sensor with an 8 kilometers radius, encompassing the entire study area. This grid was mainly designed to assess the reliability of geographical containment and the effect of sensor size in the effectiveness of the sensing process. It has allowed us to compare the effects of sensor size in the ability to capture tweets.

The second sensor set was created to generate information about specific places. We have identified several locations corresponding to parks, stadiums, airports, residential places, touristic sites and others, as listed in TABLE 1.

TABLE 1 - SENSORS FROM INDIVIDUAL PLACES

Name of the place	Type of place	Radius (Km)	Number of tweets per hour and per km²
Emirates Stadium	Stadium	0,25	29,0
Buckingham Palace and Garden	Park	0,5	26,2
Big Ben and Houses of Parliament	Landmark	0,2	304,9
Westminster Abbey	Landmark	0,2	450,6
Legoland Windsor	Park	0,64	2,5
London City Airport	Airport	2	2,1
Richmond Park	Park	2	0,2
Bristol Zoo and Westbury Park	Park	2	0,7
Hyde Park	Park	1,5	12,1
Gatwick Airport	Airport	1,5	1,0
Heathrow Airport	Airport	2	0,8
Luton Airport	Airport	2	0,5
London Southend Airport	Airport	2	0,7
Stansted Airport	Airport	2	0,2
Stamford Bridge Stadium	Stadium	0,3	104,6
Queens Park Rangers FC	Stadium	0,2	64,9
The New Den Statium	Stadium	0,2	43,2
Walthamstow Greyhound Stadium	Stadium	0,2	36,0
Arsenal Stadium	Stadium	0,2	77,7
Lloyd Park	Park	0,2	28,7
City Hall (London)	Landmark	0,2	166,9
London Business school	Landmark	0,2	100,3
British Museum	Museum	0,2	147,8
St Paul's Cathedral	Landmark	0,15	110,0
The Bank of England	Landmark	0,15	111,5
The Tower of London	Landmark	0,2	44,6
Tower's Bridge	Landmark	0,3	19,1
St Mary Axe	Landmark	0,5	12,9
The O2 (Millennium Dome)	Landmark	0,25	58,8
The SIS Building AKA the MI6	Landmark	0,15	84,0
Natural History Museum	Museum	0,15	112,4
Royal College of Art	Landmark	0,15	61,8
Lord's Cricket Ground	Stadium	0,2	91,4
Regent's Park	Park	0,7	2,7
BBC Television Centre	Landmark	0,15	111,4
Imperial War Museum	Museum	0,15	299,8
County Hall London	Landmark	0,18	105,4
Royal Festival Hall	Landmark	0,15	134,6
Royal National Theatre	Landmark	0,14	147,5
St Pancras Station	Station	0,2	48,5
Charing Cross Railway Station	Station	0,15	142,1
Ministry of Defence HQ	Landmark	0,11	1402,1
Cardinal Place London	Landmark	0,2	52,7
The National Gallery	Landmark	0,2	95,6
Haggerston Park	Park	0,25	32,2
Kennington Park	Park	0,4	10,2
Battersea Park	Park	0,6	3,5
Burgess Park	Park	0,9	4,5
Southwark Park	Park	0,5	6,9

The size of the geographical scope of these sensors was determined according to the size of the places to which they were attached. Despite the circular nature of the sensors, we tried to define the respective location and size in a way that provided the best possible match with the boundaries of the place.

The data collection took place during a period of 14 weeks, between March and June 2010. During this period, the deployed sensors have collected approximately 2027431 tweets.

B. Results regarding viability

The first objective of the evaluation was to assess the viability of these twitter-based sensors in terms of the respective data density and granularity. There were at least two key dimensions that could be considered: the technical viability of the data collection process and the ability to generate enough data for at least some of the envisioned scenarios.

Regarding technical viability, the main risk we have identified was the possible incompatibility between this usage of twitter data and the limitations imposed by twitter's API Terms of Service. In particular, the limits on the number of requests and tweets returned per request could mean that in a sensor with very high number of tweets, some tweets could be lost between requests. Therefore, a large sensor could theoretically collect fewer tweets than a set of smaller sensors corresponding exactly to the same area.

Regarding the data density generated by sensors we were particularly concerned about the granularity of the sensors, or in other words, how small could they be and still generate meaningful data. Our measure for tweeting density was based on the number of tweets per square Km per hour.

$$\boldsymbol{Tweeting\ Density} = \frac{\#\text{tweets}}{Km^2 * \ \#hours}$$

These density numbers are included in TABLE 1, for each of the places selected for sensing, and in TABLE 2, for the sensors created as part of the sensor grid.

TABLE 2 - SENSORS FROM THE GRID

Grid Level	Radius (Km)	Number of Sensors	Number of tweets per hour and per km²
Level 1	8	1	4,6
Level 2	4	4	5,6
Level 3	2	16	5,8

The numbers show that the density of the larger sensor is smaller. This may be consistent with the observation that during the data collection, the limits imposed by the Twitter service have occasionally been reached for this larger sensor, suggesting that a reliability threshold may be close to the 8km radius. The numbers also suggest that most locations, even the smaller ones, are able to sustain a continuous and significant tweeting flow. Whether or not the density values obtained in this study would be enough, will depend very heavily on the nature of the intended applications. For applications requiring real-time short-term data input, data density will be more critical. For applications that take a more long-term perspective of the data, e.g. place characterization, even much lower data densities may be enough for meaningful usage.

C. Results regarding geographical scoping

When the harvester makes a request to the twitter API, it specifies the coordinates of the center and a radius. The tweets returned as part of this request are all tweets that have been somehow connected with a particular location. In some cases, this corresponds directly to coordinates of the location, mainly when using mobile clients. In other cases, however, this has been done indirectly by associating a specific location name to the location of the person making the tweet, and then by mapping the coordinates of that location into the area specified in the request to assess whether or not it fits inside the geographical scope of the query.

The grid sensor set was conceived to combine in a single collection process multiple granularities and facilitate the study of the effectiveness of geographical scoping in the data collection process. TABLE 3 shows the number of tweets that have been detected by the grid sensors using each of those association methods. The numbers for the smaller sensors do not consider the repeated tweets that have been sensed more than once by each of those sensors.

TABLE 3 - RESULTS FOR THE SENSORS IN THE GRID

Sensors	Total of Sensors	Radius (Km)	Geo Located Tweets	Alfa Located Tweets	Total
1	1	8	1348952	429592	1778544
1A-1D	4	4	1302840	348974	1651814
1A	1	4	442411	69100	511511
1A1-1A4	4	2	405862	54548	460410
1B	1	4	521075	150638	671713
1B1-1B4	4	2	443128	110223	553351
1C	1	4	391457	100233	491690
1C1-1C4	4	2	308793	70665	379458
1D	1	4	433922	38432	472354
1D1-1D4	4	2	376735	31898	408633

We can observe from the table how the number of tweets obtained with smaller sensors compares with the number of tweets obtained from their containing sensor. The number of tweets sensed by the larger sensor is consistently bigger, but this was expectable considering that its coverage area is also bigger than the union of the coverage areas of the contained sensors.

Regarding the use of location names, we can observe that, even though location names are a relevant part of location-based tweets, about 81% of those tweets are associated with specific coordinates. The use of location names has revealed to be less reliable. Another limitation specified in the Twitter API with the use of names is that the data collection process can only be done with a minimum 1km radius. This limit does not apply to queries with geographic coordinates. We could refine this on the SpotSensing service by using the coordinates directly, but the process would be much more complex if name locations are included, as we would have to do our own mapping from location names into geographical coordinates. These results suggest that when tweeting density is critical, both methods should be used to increase the number of tweets considered. However, if the reliability of location information is more important, then only tweets with coordinates should be considered. This has lead us to include this option in the operation parameters of the twitter usage sensor.

D. Results regarding Place characterisation

The final objective of the study was to explore to what extent the data collected at different locations could provide relevant hints about the type of place being sensed, possibly providing some perception about the different nature of those places. For this part of the study we have analyzed the data collected from the sensors deployed in the various places. TABLE 4 summarizes the aggregate results obtained for the main place categories.

TABLE 4 - RESULTS FROM THE INDIVIDUAL SENSORS

	Airport	Landmark	Museum	Park	Stadium	Station
# Unique Tweets	101555	220053	13253	244737	93392	17995
# Unique Users	8293	24869	3241	18874	6254	3656
# Sensors	6	19	3	12	7	2
# Tweets/Sensor	16926	11582	4418	20395	13342	8998
# Tweets/Km^2/hour	5	3575	560	130	447	191
# Users/Km^2	119	8979	12137	459	5647	18620
% Users with a iPhone	6%	5%	13%	8%	7%	15%
% of Dialogs	31%	29%	35%	34%	29%	33%
% of English Tweets	86%	81%	86%	82%	63%	86%
% Tweets with Emotion-Tags	9%	8%	8%	10%	11%	8%
% Geo-Tweets	78%	40%	89%	61%	22%	87%
% Retweets	13%	18%	15%	14%	11%	15%

These results show how different types of place can generate different tweeting practices and open the possibility for extending this type of process for uncovering new information about places.

V. CONCLUSIONS

The results obtained from this exploratory study suggest that the overall concept of location-based social web sensor, may indeed be a viable concept.

Regarding technical viability we have studied the effect of service access limitations, which would impose a limit on the maximum size of sensors, and data density, which would impose a minimum size for sensors. We have found no problems for the sizes used in our study: 8km radius for large sensors and 2km radius for small ones. Again, in regions with less significant twitter activity, the minimum size for a meaningful twitter sensor could have to be larger, depending on the intended applications.

These types of problems will also depend on the nature of the services. For example, for a Flickr sensor, there are fewer limitations because time is not so critical, i.e. it is possible at anytime to obtain all the previously published photos at a given location, and therefore there is no risk of losing data.

Regarding the ability of the sensors to generate meaningful data about places, we have not developed any type of inference process based on sensor data, as our goal was only on the sensor concept. Still, we have shown that there are clearly distinguishing parameters that can be sensed, and therefore there is a potential to use these sensors as a source for multiple types of reality mining processes.

A. Future Work

An obvious additional feature would be to consider the tweets content, which we have not address, except for content related meta-data such as the language. For example, the words more used in a particular place could provide an important basis for characterizing that place. Also, considering that we have not explored the real applications of this sensing system, future work should consider to what extent the abstractions provided by these sensing services could be used to support inference processes that so far have been done using ad-hoc data collection processes.

REFERENCES

[1] E. Miluzzo et al., “Sensing meets mobile social networks: the design, implementation and evaluation of the CenceMe application,” in *Proceedings of the 6th ACM conference on Embedded network sensor systems*, pp. 337-350, 2008.

[2] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, e R. A. Peterson, “People-centric urban sensing,” in *Proceedings of the 2nd annual international workshop on Wireless internet*, p. 18, 2006.

[3] F. Rojas, F. Calabrese, S. Krishnan, e C. Ratti, “Real Time Rome.”

[4] B. P. Pál e P. Scientist, “SensorPlanet: An open global research framework for mobile device centric Wireless Sensor Networks,” in *MobiSensors: NSF Workshop on Data Management for Mobile Sensor Networks, Pittsburgh*, 2007.

[5] E. Paulos e T. Jenkins, “Urban probes: encountering our emerging urban atmospheres,” in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 341-350, 2005.

[6] B. Hull et al., “CarTel: a distributed mobile sensor computing system,” in *Proceedings of the 4th international conference on Embedded networked sensor systems*, pp. 125-138, 2006.

[7] M. R. Ebling e R. Cáceres, “Gaming and Augmented Reality Come to Location-Based Services,” *IEEE Pervasive Computing*, vol. 9, n. 1, pp. 5-6, 2010.

[8] A. Popescu e G. Grefenstette, “Deducing trip related information from flickr,” in *Proceedings of the 18th international conference on World wide web - WWW '09*, p. 1183, 2009.

[9] “Tweet-o-Meter - Giving you an insight into Twitter activity from around the world!.” [Online]. Available: http://www.casa.ucl.ac.uk/tom/. [Accessed: 11-Fev-2011].

[10] T. Fujisaka, R. Lee, e K. Sumiya, “Detection of Unusually Crowded Places through Micro-Blogging Sites,” in *2010 IEEE 24th International Conference on Advanced Information Networking and Applications Workshops*, pp. 467-472, 2010.

[11] S. Abrol e L. Khan, “Tweethood: Agglomerative Clustering on Fuzzy k-Closest Friends with Variable Depth for Location Mining,” in *2010 IEEE Second International Conference on Social Computing*, pp. 153-160, 2010.