

# Reality mining in urban space

Spiros Nikolopoulos  
Information Technologies Institute  
CERTH - ITI  
Thessaloniki, Greece  
Email: nikolopo@iti.gr

Symeon Papadopoulos  
Information Technologies Institute  
CERTH - ITI  
Thessaloniki, Greece  
Email: papadop@iti.gr

Yiannis Kompatsiaris  
Information Technologies Institute  
CERTH - ITI  
Thessaloniki, Greece  
Email: ikom@iti.gr

**Abstract**—The widespread adoption of smartphones and other sensor-capable devices and the availability of infrastructure capable of storing and sharing sensor information have been the main enablers for the intriguing and ambitious concept of Reality Mining. Reality mining is concerned with the collection and analysis of sensed data pertaining to different aspects of human social behavior, allowing to observe the societal dynamics and predict what individuals will do next, model the behavior of large organizations, detect trends, spot emerging phenomena and events, mine public opinion, etc. The concept of reality mining is even more intriguing when applied in the context of urban space, since user-generated content can be analyzed to reveal the citizens' perspective on various aspects of the city reality. In this paper, after elaborating on the concept of reality mining in urban space, we present three cases where user-generated content is employed to discover the citizens' perspective about: a) the city attractions, b) the city issues/problems, and c) major events in the city.

**Keywords**-reality mining; urban space; big data;

## I. INTRODUCTION

It is only recently that sensor information has become very commonplace due to the widespread availability of smartphones and other mobile sensor-capable devices. The fact that sensor data are typically weaved into mainstream social media platforms allows their combination with other online data such as text streams and updates, photos and videos. For instance, GPS, physiological, vibration, direction and health sensors applications (e.g. Garmin Connect<sup>1</sup>, EveryTrail<sup>2</sup>) collect user-contributed sensory information in real-time or a posteriori, and allow their users to share it together with other forms of social content. This unprecedented flood of data provides us with the opportunity to make decisions based on the data itself, rather than on guesswork, or on artificially constructed models of reality. Therefore, harnessing such data carries the potential to be the main enabler of *Reality Mining* by driving nearly every aspect of our modern society, including mobile services, retail, manufacturing, financial services, life and natural sciences.

*Reality mining* is concerned with the collection and analysis of sensed data pertaining to human social behavior [1]. By collecting and aggregating the data streams generated

from smartphone installed sensors (e.g. physiological measures, GPS tracks, proximity, location, and activity information) and combining them with user-generated content (UGC), we are able to gain insight into the dynamics of both individual and group behavior [2]. Such dynamics can help us predict what a single user will do next, model the behavior of large organizations, detect trends, spot emerging phenomena and events, mine public opinion, etc. What is particularly interesting in this case is the fact that mobile and social network citizens can act as the living sensors of their city, its social aspects, its cultural events or even its mobility status by doing no more than their everyday activities. Following the principles of pervasive computing the reality mining concept promises a plentiful and particularly rich content allowing to grasp the actual footprint of a living society in almost real time.

Equipped with such powerful features, the concept of reality mining proves to be useful in various application domains in urban space including: mobility, where the GPS tracks of citizens' routes can reveal the status of traffic congestion and facilitate real-time traffic planning [3]; tourism and culture, where uploaded media can reveal "off-the-beaten-path" points of interest and events, otherwise hard to discover [4]; politics, where UGC may contain up-to-date information reflecting the popularity of election candidates and their support among populations [5]; economics, marketing and brand monitoring, where the number of related media uploaded online can reflect the number and locations of products sold in the market [6]. The existence of big data streams coupling UGC with machine-sensed data and the ability to process them in large scale make feasible the extraction of the aforementioned reality aspects.

However, in achieving the ambitious objective of reality mining, a number of challenges need to be addressed, either stemming from the requirement to capture and process the large volume of data, which are expected to grow as the number of living sensors increases, or pertaining to ethical constraints, since data privacy and protection should be seriously taken into consideration when capturing the environmental and activity context of individuals.

<sup>1</sup><http://connect.garmin.com/>

<sup>2</sup><http://www.everytrail.com/>

## II. CHALLENGES FOR REALITY MINING IN URBAN SPACE

UGC produced in urban space is complex. This is due to the variety of data that it encompasses, ranging from structured data, such as logged transactions or stored measurements, to unstructured data such as text conversations, multimedia presentations and video streams, as well as due to the speed at which it is delivered and used. In using this content, we may identify the following challenges:

**Heterogeneity and Incompleteness** that call for novel data collection, fusion and aggregation methods able to combine largely different modalities coming from different sources. Such methods are necessary for exploiting the full potential of sensor information that is distributed across many sharing platforms.

**Scale and Timeliness** that raise the need for scalable and efficient data stream mining, information extraction, indexing and retrieval approaches. Such approaches need to be able to process the vast amount of data in a time interval acceptable by the application, leading to live analytics and useful information that is usually hidden in implicit relations. In many cases, since the sensors are embedded in smartphones, efficient (pre-)processing should take place in the constrained environment of the smartphone, creating distributed computing architectures.

**Privacy** relating to the fact that given the nature of big data streams (i.e. generated by use of smartphones and shared through social networks) the vast majority of potential services involves the storage and processing of users' private information. The privacy-related challenges are magnified by the fragmentation of existing policies on handling private data and the absence of a clear European directive on this issue. As a consequence users are reluctant to hand over their personal information, while researchers face many difficulties in collecting them at adequate scales.

## III. PROPOSED APPROACH

In meeting those challenges and harnessing the power of UGC to facilitate reality mining, we have successfully integrated data mining technologies with software solutions for large-scale data storage and processing. More specifically, we have developed the necessary back-end and front-end infrastructures to: a) collect and store user-generated multimedia data, b) perform the necessary processing so as to extract and exploit the collective intelligence hidden in such data, and c) engage citizens through attractive interfaces in both mobile and web platforms.

More specifically, with respect to the heterogeneity of the collected data we have incorporated in the proposed framework novel methods [7], [8] for jointly handling images and text, which are typically the most frequent resources contributed by users. In the first case [7], the rationale is to project the features extracted from the heterogeneous resources into a new space (e.g. space of probabilities) where all data samples can be treated seamlessly. In the latter [8],

visual and textual features of user-contributed images are fused by use of a similarity graph, employed for clustering.

For coping with scalability and real-time processing requirements, we rely on the best practices for distributed processing and storage, such as the MapReduce [9] and the MPI (Message Passing Interface) paradigms, and the schema-less auto-sharding MongoDB database<sup>3</sup> for enabling flexible data modelling and horizontal scaling. Computational efficiency is also achieved at the level of image processing where highly optimized compact image descriptors are employed to power similarity search [10]. Finally, for dealing with privacy-related challenges we make sure to inform users about the potential usage of their data, as well as to obtain their consent before collecting any information. Moreover, when using shared data (e.g. Flickr images), appropriate licenses are selected that permit content reuse and redistribution (e.g. Creative Commons licenses<sup>4</sup>).

Based on the above technologies, we present three cases, where reality mining has been applied on UGC to mine certain aspects of urban reality, namely, city attractions, city issues/problems, and major events that take place in the city.

## IV. CASE STUDIES

### A. Mining the city attractions

The first case study involves an application that applies multimedia mining techniques on large amounts of UGC to discover interesting areas, places of interest (POIs) and events in a city<sup>5</sup>. The application employs a two-level clustering scheme on very large collections of geotagged images collected from Flickr. At the first level, images are automatically clustered by location into areas or districts of interest (e.g. Downing street in the city of London). Then, at a finer clustering level, images of the same district are clustered around landmarks and events, and titles are extracted for each cluster based on the metadata of the associated images. The entire process takes place in a fully automatic way enabling the creation of online city guides at a global scale. More details about the techniques behind this application can be found in [8]. A set of 40 cities are currently included in the application. A snapshot from the mobile app UI is presented in Figure 1.

### B. Mining the city issues/problems

In modern societies there is a growing requirement for public administrations to directly communicate with their citizens, view the existing problems from their perspective and react to their needs. To this end, smartphones have turned into a valuable instrument that, apart from being a source of rich information, is also an integral part of our daily activities. Based on the above and driven by the

<sup>3</sup><http://www.mongodb.org/>

<sup>4</sup><http://creativecommons.org/>

<sup>5</sup><http://www.clusttour.gr/>

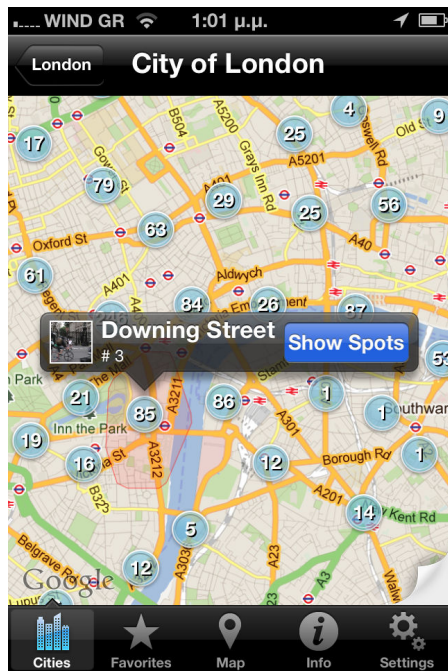


Figure 1. Spots identified in a city




Figure 2. Map-based view of all issues in a city

intriguing concept that citizens can act as living sensors of their city, our second use case consists of a user-friendly mobile application that enables them to directly report administration problems (or suggestions) about their city, such as potholes, illegal trash dumping, faulty street lights, broken tiles on sidewalks, and illegal advertising boards. The issues submitted by citizens are subsequently transmitted, through the back-end infrastructure, to the appropriate office in public administration so as to schedule their settlement.

The main functionalities offered by the application<sup>6</sup> include: a) a map-based and a list-based view of the reported issues (Figure 2) that are also supported by filtering mechanisms, allowing the user to quickly obtain an overview of the existing problems, b) a “new issue” reporting functionality that allows the user to easily describe, locate, visualize and submit an issue, c) a mechanism for all registered users to comment on reported issues or give their positive vote, and e) a feedback mechanism for informing the citizens about the current progress of their issues.

The mobile application communicates with an integrated back-end environment for managing the incoming issues. This environment provides the necessary interfaces for obtaining an overview of the reported issues and understanding the citizens’ perspective. Finally, reality mining is also facilitated through the generation of statistically aggregated information, in time and space, that allows the city officials to assess the overall performance of the municipality and plan for future resources.

<sup>6</sup><http://mklab-iti.github.io/ImproveMyCity-Mobile/>

### C. Reality mining about major events in a city

The third case study pertains to a platform that enhances the experience of people attending large events, such as festivals, by enabling the extraction of rich insights with respect to the impact of such events to their attendants. More specifically, the proposed platform includes a mobile app targeted at the event of interest and a server-side data mining component that processes the streams of social media focused on the event and the inputs of event attendants contributed through the app. The mobile app facilitates rapid lookup of the festival program and enables users to create personalized views of the program. More importantly, the app enables users to give feedback on the event, e.g. in the case of a film festival, users can rate and comment on the screened films (see Figure 3). The server-side component aggregates user ratings and bookmarks to derive popularity rankings of films. At the same time, it monitors the Twitter stream for messages related to the festival and employs text mining to automatically match them to films, discover emerging topics, and detect their sentiment (positive/negative/neutral). More technical details on the implementations of these approaches and their assessment are available in [11].

By combining and mining these two sources (mobile app inputs and Twitter stream), the developed event platform helps event organizers to obtain a succinct and informative view of the event in question. For instance, in the case of the recent Thessaloniki International Film Festival (TIFF53) Figure 4 provides a summary of the Twitter activity in

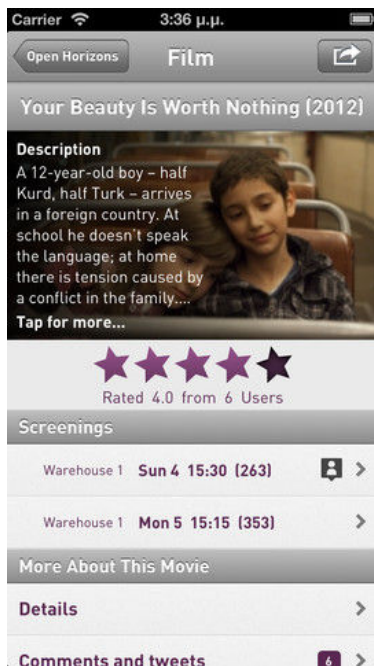


Figure 3. Mobile app for festivals




Figure 4. Twitter activity during TIFF53

tandem with the most important topics detected by mining online discussions. In addition, it can provide lists of top ranked films, topics, shared URLs, and media content. In some cases, the platform could help organizers identify potential problems (e.g. attendants' complaints about services offered in the event vicinity) and understand what people like in the event.

## V. CONCLUSION

Reality mining is an intriguing and ambitious concept that has only started to attract the attention of researchers due to the abundant availability of user generated sensor data and content. Researchers are intrigued by the opportunities and potential of real-time urban data, including the capturing of city landmarks, issues and the citizens' experiences in

the context of large events. Nevertheless, in order for such data to lead to meaningful conclusions, they need to be captured and analyzed in large scale, giving rise to certain challenges associated with the nature of such data. The applications presented in this paper demonstrate how data mining technologies can be combined with large-scale multimedia storage and processing approaches to mine large amounts of UGC and gain insights into different aspects of urban reality.

## ACKNOWLEDGMENT

This work was supported by the FP7-288815 project Live+Gov and the FP7-287975 project SocialSensor.

## REFERENCES

- [1] N. Eagle and A. (Sandy) Pentland, "Reality mining: sensing complex social systems," *Personal Ubiquitous Comput.*, vol. 10, no. 4, pp. 255–268, Mar. 2006.
- [2] A. Thiagarajan, J. Biagioni, T. Gerlich, and J. Eriksson, "Cooperative transit tracking using smart-phones," in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '10. New York, NY, USA: ACM, 2010, pp. 85–98. [Online]. Available: <http://doi.acm.org/10.1145/1869983.1869993>
- [3] Detection of traffic congestion using gps tracks. [Online]. Available: <http://googleblog.blogspot.gr/2009/08/bright-side-of-sitting-in-traffic.html>
- [4] C. Zigkolis, S. Papadopoulos, Y. Kompatsiaris, and A. Vakali, "Detecting the long-tail of points of interest in tagged photo collections," in *Content-Based Multimedia Indexing (CBMI), 2011 9th International Workshop on*, 2011, pp. 235–240.
- [5] Record and analyze social media for predictive purposes. [Online]. Available: <https://www.recordedfuture.com/about/>
- [6] Social media monitoring and engagement platform. [Online]. Available: <http://brandmonitor.thismoment.com/>
- [7] S. Nikolopoulos, S. Zafeiriou, I. Patras, and I. Kompatsiaris, "High order pls for indexing tagged images," *Signal Process.*, vol. 93, no. 8, pp. 2212–2228, Aug. 2013.
- [8] S. Papadopoulos, C. Zigkolis, Y. Kompatsiaris, and A. Vakali, "Cluster-based landmark and event detection for tagged photo collections," *IEEE MultiMedia*, vol. 18, no. 1, pp. 52–63, Jan. 2011.
- [9] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," *Commun. ACM*, vol. 51, no. 1, pp. 107–113, Jan. 2008. [Online]. Available: <http://doi.acm.org/10.1145/1327452.1327492>
- [10] E. Spyromitros-Xioufis, S. Papadopoulos, I. Kompatsiaris, G. Tsoumakas, and I. P. Vlahavas, "An empirical study on the combination of surf features with vlad vectors for image search," in *WIAMIS*, 2012, pp. 1–4.
- [11] E. Schinas, S. Papadopoulos, S. Diplaris, Y. Kompatsiaris, Y. Mass, J. Herzig, and L. Boutakidis, "Eventsense: Capturing the pulse of large-scale events by mining social media streams," in *17th Panhellenic Conference on Informatics (PCI)*, 2013.