Challenges in Recommending Venues within Smart Cities

Romain Deveaud, M-Dyaa Albakour, Craig Macdonald, and Iadh Ounis
University of Glasgow, UK
firstame.lastname@glasgow.ac.uk

ABSTRACT

Recommending venues to a user within a city is a task that has emerged recently with the growing interest in location-based information access. However, the current applications for this task only use the limited and private data gathered by Location-based Social Networks (LBSNs) such as Foursquare or Google Places. In this position paper, we discuss the research opportunities that can arise with the use of the digital infrastructure of a smart city, and how the venue recommendation applications can benefit from this infrastructure. We focus on the potential applications of social and physical sensors for improving the quality of the recommendations, and highlight the challenges in evaluating such recommendations.

Keywords

social and physical sensors, location-based social networks, contextual evaluation

1. INTRODUCTION

With the wide use of Internet-connected and sensor-enabled smartphones, users can now share their location while performing online tasks, such as searching for information. The collected location data allows to model the behaviour of the users, such as their travel preferences, thereby enabling the emergence of new tasks such as the recommendation of venues that might be of interest to the user. This task encompasses a wide variety of sub-tasks, ranging from the recommendation of venues that the user have never visited before [4] to the prediction of the next location of the user [3].

Currently, the existing venue recommendation systems heavily rely on data extracted from Location-based Social Networks (LBSNs) [7], such as Foursquare[1], Yelp[2], or Google Places[3]. In these LBSNs, users can broadcast their location to their friends (or to other users), and rate and comment on venues. In the literature, the standard approach for recommending venues is to identify users that are similar to the current user, and to recommend venues that these similar users have rated highly [8, 9].

Since the aforementioned approaches rely on the profiles of the users, which are composed of private data, the only existing applications, thus far, are industrial and are provided by the LBSNs. Such applications examine the history of the visits of the users, and recommend venues that might be of interest to them. However, the recommendations only rely on a few signals (i.e. the derived preferences of a user through the venues she/he visited before) to estimate the relevance of venues. The privacy of the user data also raises questions about the generalisation and the reusability of such venue recommendation approaches. We argue that the digital infrastructure provided by smart cities can overcome these problems. In this position paper, we discuss the new possibilities and the underlying challenges of using new types of data for recommending venues. In Section 2, we discuss the use of different signals and indicators found in smart cities deployments that might improve the representation of venues, while we focus in Section 3 on the problems related to the evaluation of such recommendations.

2. ON THE USE OF SENSORS

One of the main characteristic of smart cities is the abundance of sensors connected to the Internet of Things [5], which allow to gather information on what is currently happening in the city. While we may think primarily of physical sensors (e.g. CCTV cameras), we may also consider social sensors (e.g. Twitter [1]). Currently, only social sensors are used to perform venues recommendation. Indeed, the collaborative filtering recommendation methods that are typically used in the literature do not use other indicators than the profiles of the users. Moreover, the preferences inferred from these profiles cannot be shared amongst different LBSNs [4]. Combining data from different social sensors could improve the representation of venues, and eventually overcome the sparsity problem that can arise for venues that have few associated ratings and comments. Since some tweets are geolocated, it is possible to analyse the sentiments [6] expressed in the tweets that were emitted at the locations corresponding to the venues, and use these sentiments as a sensor of the quality of the venue.

In addition to a combination of different social sensors, using physical sensors allows to derive further information about the venues. For example, CCTV cameras and microphones, along with some audio/video processing, can be used to estimate crowd densities, thus helping to identify in real-time popular areas (including the corresponding venues) that might be of interest to a user. Such physical observations may also help to better model the context of the users by identifying their behaviours (e.g. a venue recommendation system should not suggest outdoor venues to someone who appears to be feeling cold while walking in the street). The GPS sensors on public transports can also help to detect traffic jams (and even predict them, with sufficient training

[1] http://foursquare.com/
[2] http://yelp.com
[3] http://maps.google.com/

ECIR'14 Information Access in Smart Cities Workshop (i-ASC 2014). April 13, 2014, Amsterdam, the Netherlands.

data), which might prevent users to promptly reach certain venues. Environment sensors (e.g. rain-related) can also help to determine areas that might not be suitable to recommend (e.g. outdoor venues), or that might be difficult to reach for certain types of users (e.g. elderly persons).

We argue that having rich representations of real-world entities, such as venues, is essential for providing high-quality and accurate recommendations to the users. While social sensors are essential sources of subjectivity (e.g. opinions/ratings about a venue), physical sensors can provide valuable additional objective indicators that can help to identify the context of the venue.

However, there is a need for new effective methods that can interpret all of the raw data extracted from these physical sensors, in order to generate useful information. There is also a need for agreed standards for storing this data (e.g. RDF). One can imagine that all sensors could feed a dedicated knowledge base of the smart city that holds all the records of the different entities of the city. Internet-connected sensors can then update the attributes of the entities in real-time. Such knowledge bases will need however to evolve from a static representation of the information to a time-aware representation, allowing to track the evolution of each entity's attributes and to eventually forecast them. Finally, it is of note that special considerations should be paid to the privacy and ethical issues arising from the storage of such data. Such issues still need to be explored so as to ensure an adequate balance between added-value and privacy.

3. EVALUATING RECOMMENDATIONS

Although a variety of new ideas can be imagined to improve the recommendation of venues, they need to be tested and properly evaluated. However, we face here again a challenge that is related to the nature of the task itself. Indeed, a venue recommendation is contextual to (among other parameters) the time of the day, the current location of the user, and her/his preferences. The TREC Contextual Suggestion track [2] explored such research questions, aiming to develop a test collection that can support the venue recommendation task. In its first year (2012), the track explored the relevance of recommendations with respect to a description of the venue generated by the systems, the geographical relevance, the adequacy of the website linked to the venue, and the temporal relevance. Several evaluation measures combining these aspects were also proposed. In its second year (2013) however, the time aspect was dropped, thus highlighting the difficulties in evaluating venue recommendations with respect to all the contextual parameters at stake. Moreover, other parameters could be considered. We raised the issue of privacy in the previous section; some users might want to pay the price of giving their personal information in order to benefit from highly relevant recommendations, while others might prefer to receive only "good" enough recommendations by sharing less personal information. Such evaluation frameworks will then need to take a wide range of parameters into account in order to accurately estimate the relevance of recommendations to users.

Building an entirely reusable test collection for evaluating venue recommendations is also a challenge. Indeed, we argued that the relevance of a venue recommendation is highly contextual and can change depending on a wide range of parameters, which may be difficult to reproduce in a controlled setting. While simplified settings are required to understand how the proposed systems perform, they do not necessarily reflect real use cases. Using a smartphone application, and analysing the feedback provided by the users, is a possible way to evaluate the quality of the recommendations. This feedback can be of two different types: explicit or implicit. An explicit feedback takes the form of a questionnaire, asking the user if the recommendation is interesting. This questionnaire can also be presented before and after the visit, in order to analyse if the experience actually changed the opinion of the user about the venue. An implicit feedback can be gathered from the GPS data of the smartphone: if users actually went to a venue that the application recommended, then they must have found the recommendation interesting given their context. The implicit evaluation techniques used by commercial search engines, such as A/B testing, could also be used as an implicit feedback.

4. CONCLUSION

The task of recommending venues to users is an emerging task that presents many challenges. In this position paper, we argued that this task can benefit from the digital infrastructure deployed by smart cities. More particularly, the use of physical sensors offers advantages that have currently remained unexplored. We have also argued that sensing the behaviours of users when they interact with their mobile devices is also a promising direction for accurately evaluating the quality of recommendations.

5. ACKNOWLEDGMENTS

This work has been carried out in the scope of the EC co-funded project SMART (FP7-287583).

6. REFERENCES

[1] M.-D. Albakour, C. Macdonald, and I. Ounis. Identifying Local Events by Using Microblogs As Social Sensors. In *Proc. of OAIR*, 2013.

[2] A. Dean-Hall, C. L. Clarke, J. Kamps, P. Thomas, N. Simone, and E. M. Voorhees. Overview of the TREC 2013 Contextual Suggestion track. In *Proc. of TREC*, 2013.

[3] A. Monreale, F. Pinelli, R. Trasarti, and F. Giannotti. WhereNext: A Location Predictor on Trajectory Pattern Mining. In *Proc. of KDD*, 2009.

[4] A. Noulas, S. Scellato, N. Lathia, and C. Mascolo. A Random Walk Around the City: New Venue Recommendation in Location-Based Social Networks. In *Proc. of SOCIALCOM-PASSAT*, 2012.

[5] J. Soldatos, M. Draief, C. Macdonald, and I. Ounis. Multimedia Search over Integrated Social and Sensor Networks. In *Proc. of WWW*, 2012.

[6] M. Thelwall, K. Buckley, G. Paltoglou, D. Cai, and A. Kappas. Sentiment in Short Strength Detection Informal Text. *J. Am. Soc. Inf. Sci. Technol.*, 61(12):2544–2558, 2010.

[7] M. Ye, P. Yin, and W.-C. Lee. Location Recommendation for Location-based Social Networks. In *Proc. of SIGSPATIAL GIS*, 2010.

[8] M. Ye, P. Yin, W.-C. Lee, and D.-L. Lee. Exploiting Geographical Influence for Collaborative Point-of-interest Recommendation. In *Proc. of SIGIR*, 2011.

[9] Q. Yuan, G. Cong, Z. Ma, A. Sun, and N. Magnenat-Thalmann. Time-aware Point-of-interest Recommendation. In *Proc. of SIGIR*, 2013.