Out of the Bubble - Serendipitous Event Recommendations at an Urban Culture Festival

Andreas Forsblom[1,2], **Petteri Nurmi**[1,2], **Pirkka Åman**[2], **Lassi Liikkanen**[2]
[1] Helsinki Institute for Information Technology HIIT
PO Box 68, FI-00014 University of Helsinki, Finland
firstname.lastname@cs.helsinki.fi
[2] Helsinki Institute for Information Technology HIIT
PO Box 19800, FI-00076 Aalto University, Finland
firstname.lastname@hiit.fi

ABSTRACT
Advances in positioning technologies have resulted in a surge of location-based recommendation systems for mobile devices. A central challenge in these systems is to avoid the so-called filter bubble effect, i.e., that people are not only exposed to information that is in line with their personal ecosystem, but that they can also discover novel and otherwise interesting content. We present results from a field study of a mobile recommendation system that has been aimed to support serendipitous discovery of events at an urban culture festival. Results from the study indicate that suitably designed recommendations together with access to relevant external information sources can lead to serendipitous discovery of new content, such as new artists, bands or individual songs. Our results also indicate that proximity has little effect on the effectiveness of serendipitous recommendations.

Author Keywords
Mobile Recommendations, Serendipity, Urban Computing

ACM Classification Keywords
H.5.m. [Information Interfaces and Presentation (e.g. HCI)]: Miscellaneous; H.4.2. [Information Systems]: Information Systems Applications

General Terms
Human Factors, Experimentation

INTRODUCTION
Increased sensing capabilities of smartphones combined with advances in positioning technologies have resulted in a widespread interest in mobile recommender systems that tailor their content based on the user's current location, personal profile, activity or other situational factors; see, e.g., [1, 4, 7, 8]. A central challenge in these systems is how to avoid the so-called *filter bubble*[1], i.e., that people are only exposed to content that is in line with their personal ecosystem. Another major challenge is avoiding the *popularity bias*, i.e., that recommendations and consumption tend to concentrate on popular items instead of enabling people to easily find novel and serendipitous content [3]. To explore how mobile recommender systems can support serendipitous discovery in an urban context, we have developed Sounds of Helsinki (SoH), a novel mobile recommendation system. To evaluate the effectiveness of SoH in supporting serendipitous discovery, we have conducted a field study with 15 participants during an urban culture festival in Helsinki, Finland. In the study, SoH was used to recommend cultural events in the urban environment. The recommendations were augmented with additional media whenever possible. The results of the study demonstrate that suitably designed recommendations can lead to serendipitous discovery of new content, such as artists, bands or individual songs. As part of the study we also evaluated whether proximity to recommended events had an influence on the effectiveness or perceived relevance of the presented recommendations.

SOUNDS OF HELSINKI
Sounds of Helsinki is a mobile recommendation system for supporting serendipitous discovery of events in an urban environment. Sounds of Helsinki has been implemented following a client-server architecture. The server-side is responsible for generating event recommendations, managing event information and maintaining information about the location of the user. The server-side also acts as repository for any additional media that is available about the performers related to the events. The client-side of the system is web-based and operates in the browser of the mobile phone. The client-side additionally consists of a background script that periodically sends the GPS location of the user to the server. The script is also responsible for triggering notifications about new recommendations whenever they become available. The notifications are given using vibro-tactile feedback and playing an

[1] http://www.ted.com/talks/eli_pariser_beware_online_filter_bubbles.html

IUI'12, February 14–17, 2012, Lisbon, Portugal.

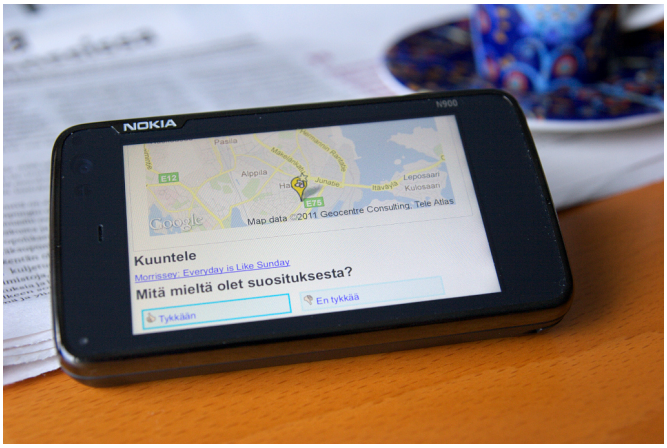

Figure 1. The user interface displaying part of a recommendation on a Nokia N900.

audio alert. To avoid unwanted disruptions, the notifications conform with the current profile of the client device.

The web based user interface is designed to run on a Nokia N900 smartphone, which has a 3.5" touch screen. The interface has two main views, a recommendation listing and an event view. The recommendation listing shows all of the user's recommendations in reverse chronological order, with previously unseen ones marked as new. This view is opened automatically by the client when a new recommendation becomes available. This view can also be opened manually at any time using a shortcut on the phone's desktop. The recommendations can be viewed by tapping on the corresponding item on the listing, which leads to the event view. The event view shows the name of the event and the performer(s), the address of the venue where the event takes place and distance from the user's current position, static maps centered on the user and on the event, and an overview map displaying both; see Fig. 1. In addition, the time and date of the event is shown. For recurring events, all future showings or performances are also listed. If the event is sold out, this is also indicated. For events that we were able to locate suitable media for, the event view enabled the user to listen to music samples of the artist, or to view embedded YouTube videos of the artist.

Sounds of Helsinki currently integrates two different recommendation techniques. Both recommendation techniques restrict the recommendations to events that take place during the current day and which have not yet started. We refer to the events satisfying these criteria as the *candidate pool*. Both recommendation techniques assign a weight w_i for each of the events i in the candidate pool C and select the event to recommend through weighted random sampling. The first technique, `RANDOM`, assigns uniform weights for all events in the candidate pool. The second technique, `LOCATION-BASED`, assigns the weight based on proximity to the event. Specifically, if d_i denotes the distance in kilometers to the event i from the user's current location, the `LOCATION-BASED` recommender uses $\exp(-d_i)$ as the weight of event i. As sold-out events would be associated with information about artists and potentially also with media, we did not filter out sold-out events from the candidate pool before generating recommendations.

FIELD STUDY

To evaluate the effectiveness of Sounds of Helsinki in supporting serendipitous discovery, we have conducted a field study during the Helsinki Festival 2011, a yearly cultural festival taking place in Helsinki, Finland. The duration of the festival is approximately two weeks, during which numerous cultural events, such as music, dance, cinema and theater performances take place at different locations throughout the city. For conducting the user evaluation, we recruited two groups of people that described themselves as active listeners of mobile music and avid event-goers. The participants were aged between 24 and 45. The first group ($N = 8$) participated in the study for two days, while the second group ($N = 7$) participated for three days that did not overlap with the days of the first group. The participants were provided with Nokia N900 smart phones equipped with prepaid SIM cards for the duration of the study, and they were rewarded with tickets to one of two selected events during the festival. The reward events were excluded from being recommended by our system.

To avoid disrupting participants during night or working hours, recommendations were only generated between 2 p.m. and 10 p.m. To preserve battery of the client device, GPS location updates were sent to the server at five minute intervals. The time between successive recommendations was randomly sampled between 40 and 80 minutes. The algorithm that was used to generate recommendations was chosen randomly each time a new recommendation was made. For each recommendation, the participant had the option of rating the recommendation using *thumbs up* and *thumbs down* buttons.

To supplement the thumbs up/down ratings, at the end of the field study we collected post-hoc survey responses. Specifically, for each participant, we replayed the event recommendations that were shown to him/her during the study. For each recommendation, the participant was asked to respond to a six item questionnaire that measured whether the participant saw the recommendation, how good and serendipitous the recommendation was, and whether the recommendation influenced the participant's behavior or not. The question items that were included were adapted from the user-centric recommendation evaluation framework proposed by Pu et al. [6]; see Table 1 for the items that were included in the questionnaire. Responses for all items were elicited using a 5-point Likert scale that was anchored at "completely disagree" (=1) and "completely agree" (=5). After completing the questionnaire for each event recommendation, the participant was asked to complete a ten item end-questionnaire measuring the overall usefulness of the system and its capability to support serendipity; see Table 2.

RESULTS

Events and Recommendations

In total, 241 recommendations were made in the field study. Most of the recommendations (212, 88%) were viewed by

the participants, which indicates that the system was effective in increasing awareness of available events. In total, recommendations were made for 68 different events. Out of 241 recommendations, 95 (39%) had either streaming music or a YouTube video associated with them. The media was consumed at least once in 50 of the 95 cases (53%), but more than once only in 5 cases. Of the recommendations that were viewed, 74 (35%) were rated as good (thumbs up), 46 (22%) as poor (thumbs down), and 92 recommendations (43%) were not rated. No statistically significant differences were observed in the ratings between the RANDOM or the LOCATION-BASED recommendations. The median distance to events recommended using the RANDOM technique was 2.9 kilometers from the user's current location. For the LOCATION-BASED recommendations the median distance was 2.0 kilometers.

Influence of Recommendations

Table 1 summarizes responses elicited to questions measuring the relevance and goodness of individual event recommendations. We used non-parametric ANOVA to compare the responses to the recommendations generated with the RANDOM method and the LOCATION-BASED method. The comparison revealed no statistically significant differences between the two recommendation methods. We separately evaluated whether distance to recommended events influenced any of the responses. This analysis was conducted using one sided t-tests assuming unequal variance. Before analysis, the distance values were normalized by applying a logarithmic function on them. The analysis did not return any statistically significant effects, which suggests that proximity had no influence on perceived relevance or effectiveness of recommendations in our study.

The responses to the questionnaire indicate, first of all, that most of the recommendations were for artists/performers that were previously unknown to the participant and for events that the participant had not heard of before. Together with the high consumption rate of additional media, the results thus indicate that SoH was effective in supporting serendipity. However, the responses also indicate that the event recommendations were considered neither interesting nor uninteresting, implying a trade-off between relevance of recommendations and support for serendipity.

The recommendations that were presented were considered neither good or bad in terms of their timeliness. The responses also suggest that the recommendations had little impact on the behavior of the participants. Comments of the participants suggest that one of the main reasons for the limited impact of the recommendations on their behavior was related to pre-planning of the day's schedule before being exposed to the event recommendations. Some of the recommended events that the participants considered interesting had already been sold out which prevented them from having an influence on the participant's behavior.

Table 2 summarizes the responses to the questionnaire measuring the overall goodness of the system. The responses indicate, first of all, that, while individual event recommendations were not considered particularly relevant, overall the participants were satisfied with the system and the support it provided for determining which events to attend. Similarly to the responses to the individual event recommendations, the participants did not consider the system to have a significant influence on their behavior.

DISCUSSION AND CONCLUSIONS

We have presented Sounds of Helsinki, a mobile recommendation system for supporting serendipitous discovery of events in an urban environment. Results from a field study with 15 participants, conducted during an urban culture festival, indicate that SoH is effective in supporting serendipity. However, the results also indicate that the recommendations were not considered particularly interesting or relevant. While personalization techniques could be used to improve the interestingness of recommendations, this would easily decrease the effectiveness of SoH in supporting serendipity. A potential solution would be to use personalization techniques not only for determining what to recommend, but to determine when the user is more likely to be interested in novel and serendipitous content.

The results of the study also indicated that proximity to events had no influence on the perceived goodness of the recommendations. The city of Helsinki, where the study was conducted, has a good public transportation system and many central culture venues are within walking distance from each other in the city centre. For a person located outside the city center, this means that it will take him or her approximately equally long to get to any venue. On the other hand, for someone already in the city centre, most culture venues can be reached within 15 minutes by foot or by public transportation. Given this, the effort required to get to an event is likely to be similar for events that are recommended using either of the recommendation techniques. The similar rate of positive and negative responses for the two recommendation techniques further supports this interpretation. This finding is in line with results from studies in location-based personal information management [5] and mobile marketing [2] which have shown that user preferences regarding location are complex, extending beyond geographic proximity and depending on a multitude of factors ranging from user plans to movement routines and social geography of an area.

Acknowledgments

The authors wish to thank Helsinki Festival for providing the reward tickets. This work was supported by Aalto MediaFactory and EIT ICT Labs.

REFERENCES

1. Bellotti, V., Begole, B., Chi, E. H., Ducheneaut, N., Fang, J., Isaacs, E., King, T., Newman, M. W., Partridge, K., Price, B., Rasmussen, P., Roberts, M., Schiano, D. J., and Walendowski, A. Activity-based serendipitous recommendations with the Magitti mobile leisure guide. In *Proceeding of the 26th annual SIGCHI conference on Human factors in computing systems*, ACM (2008), 1157–1166.
2. Drossos, D., Giaglis, G. M., Lekakos, G., Kokkinaki, F., and Stavraki, M. G. Determinants of effective SMS

	RANDOM	LOCATION-BASED
	Median (IQR)	Median (IQR)
I am sure that I saw this recommendation.	5 (2)	5 (1)
The recommendation was interesting.	3 (2)	3 (2)
I was already familiar with the artist.	1 (3)	1 (2)
I had already heard about the recommended event.	1 (1)	1 (1)
The recommendation was presented at a suitable time.	3 (2)	3 (3)
The recommendation affected my behaviour.	1 (1)	1 (1)

Table 1. Summary of per-recommendation questionnaire results. The questions have been translated from Finnish.

	Median (IQR)
I enjoyed the [listening to or watching] the artists/bands of the recommended events (through media samples or attending the events).	3 (1.0)
The recommendations helped me find new artists.	3 (1.5)
The artists/bands related to the recommended events were not previously familiar to me.	3 (2.0)
I found the recommended events interesting.	4 (0.5)
The recommended events did not interest me.	2 (1.0)
The recommended events were similar to each other.	2 (0.5)
The recommendations were timely	4 (1.0)
The recommendations affected my decisions to attend events.	2 (2.0)
When I use the service, I feel that it supports me in trying to find content or events that I like.	4 (1.0)
Overall, I am satisfied with the recommendations made by the system.	4 (1.0)

Table 2. Summary of end-questionnaire results. The questions have been translated from Finnish.

advertising: An experimental study. *Journal of Interactive Advertising 7*, 2 (2007), 16–27.

3. Fleder, D. M., and Hosanagar, K. Recommender systems and their impact on sales diversity. In *Proceedings of the 8th ACM conference on Electronic commerce*, ACM (2007), 192–199.
4. Krüger, A., Baus, J., Heckmann, D., Kruppa, M., and Wasinger, R. Adaptive mobile guides. In *The Adaptive Web*, P. Brusilovsky, A. Kobsa, and W. Nejdl, Eds., vol. 4321. Springer, 2007, 521 – 549.
5. Ludford, P. J., Frankowski, D., Reily, K., and Terveen, L. Because I carry my cellphone anyway: Functional Location-Based Reminder Applications. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI'06)* (2006).
6. Pu, P., Chen, L., and Hu, R. A user-centric evaluation framework for recommender systems. In *Proceedings of the 2011 ACM Conference on Recommender Systems (RecSys)*, ACM (2011), 157–164.
7. Tintarev, N., Flores, A., and Amatrain, X. Off the beaten track - a mobile field study exploring the long tail of tourist recommendations. In *Proceedings of the 12th International Symposium on Human-Computer Interaction with Mobile Devices and Services (MobileHCI)* (2010).
8. Zheng, V. W., Zheng, Y., Xie, X., and Yang, Q. Collaborative location and activity recommendations with GPS history data. In *Proceedings of the 19th international conference on World wide web*, ACM (2010), 1029–1038.