

Map-based Interaction with a Conversational Mobile Recommender System

Olga Averjanova, Francesco Ricci, and Quang Nhat Nguyen
Free University of Bolzano
{oaverjanova, fricci, qnhatnguyen}@unibz.it

Abstract

Recommender systems are information search and decision support tools used when there is an overwhelming set of options to consider or when the user lacks the domain-specific knowledge necessary to take autonomous decisions. They provide users with personalized recommendations adapted to their needs and preferences in a particular usage context. In this paper, we present an approach for integrating recommendation and electronic map technologies to build a map-based conversational mobile recommender system that can effectively and intuitively support users in finding their desired products and services. The results of our real-user study show that integrating map-based visualization and interaction in mobile recommender systems improves the system recommendation effectiveness and increases the user satisfaction.

1. Introduction

The advent of mobile devices (e.g., PDAs and mobile phones), wireless communication technologies (e.g., wireless LAN and UMTS), and position detection techniques (e.g., RFID beacon-based and GPS), has made possible the development and commercialization of a large number of location-based services (LBSs) [13, 6]. There exists now a wide range of different LBSs, including emergency services, information services, navigation support services, etc. For example, mobile users can access local tourist information services providing information about nearby points of interests [9], such as pubs and restaurants [3], or get routing guidance from their position to a target location [9]. In many LBS systems, maps are used to visualize various kinds of information related to the points of interest (e.g., opening hours) and their spatial relations.

As the amount of online information and services continuously increase, it becomes more and more

difficult for users to find the right information required to accomplish a particular task (e.g., travel planning). Moreover, given the limitations of mobile devices, such as small screen and limited computing power, displaying in an electronic map a large number of information objects, and related information, is computationally expensive and not effective. Hence, LBS systems should employ filtering mechanisms to reduce the amount of data displayed on electronic maps.

Recommender systems [11, 1] intend to address that problem, which is also called “information overload”, by suggesting users the products and services best matching their needs and preferences. Recommender systems are useful not only when users are overwhelmed by a large number of options to consider but also when they do not have enough domain-specific knowledge to take their decision [1]. Notwithstanding the stated benefits of map-based interfaces and recommendation technologies, the integration of these two technologies and the evaluation of the usability and effectiveness improvements, which they can bring to a mobile information system, have not been studied yet.

In this paper we discuss the design and evaluation of a map-based mobile recommender system that, as we shall illustrate, can effectively assist users in finding their desired products and services. In particular, we empirically evaluate if such an integrated approach can improve a mobile recommender system previously developed, called MobyRek [12]. MobyRek is a critique-based conversational recommender system that supports mobile users in finding their desired travel products (restaurants). The results of some previous experiments, consisting of a live-user evaluation [12] and a number of simulation tests [7, 8], showed that the MobyRek recommendation methodology is effective. Nevertheless, we hypothesized that some limitations of MobyRek’s user interface (discussed in Section 3) may cause difficulties and inconveniences for users in their interactions with the system.

In this paper, we identify these usability limitations of MobyRek, which are general limitations of mobile recommender systems, and we present an extended and improved version of the system called MapMobyRek. MapMobyRek uses maps as the main interface for items access and information display and provides new decision-support functions based on the map. The design of the map-based interface of MapMobyRek focused on the following user functions:

- to enter the search query by specifying preferences for item features,
- to see the system's recommendations on the map,
- to recognize immediately the differences between good and weak recommendations,
- to compare two selected recommendations,
- to input critiques to the recommended items,
- to see on the map how the expressed critique influences the system's recommendations, and
- to select the best items.

We present in this paper a usability test aimed at comparing the two systems, MobyRek and MapMobyRek, with respect to functionality, efficiency, and convenience. The test's results show that MapMobyRek does improve MobyRek with respect to all these three dimensions.

The paper is organized as follows. In Section 2 we survey related work. In Section 3 we present our recommendation methodology, discuss the limitations of MobyRek's user interface, and present MapMobyRek – the extended system. In Section 4 we state our research hypotheses and present the test procedure. In Section 5 we report on and discuss the test results. Finally, in Section 6 we give our conclusions and indicate some future work.

2. Related work

There have been research works dealing either with map-based mobile services [2, 5] or with recommender systems [1], separately. However, here we focus on related works addressing the intersection of map-based mobile applications and recommender systems.

CityGuide [3] is designed for PalmOS devices to help tourists find attractions (such as restaurants) around a city. This recommender system uses the constraint-based filtering approach to control which attractions are shown on the map. In particular, the user, through the system's map interface, is asked to specify constraints on attraction type, restaurant cuisine and price. The user's indicated constraints are then used by the system to compute the attractions shown on the map.

Dynamic Tour Guide (DTG) [14] is a mobile tour guide system that helps travelers in discovering a destination. Given a user's request, the system computes a personalized tour in the city by asking the user her interests and the time she would like to spend for the visit. In DTG a tour is composed by a set of points of interest (i.e., restaurants, attractions and events). The system then visualizes the recommended tour on the map and provides some audio guide information. The system computes and visualizes just one tour per each user request, instead of providing a list of recommended tours.

Mobile Piste Map [4] is designed to support skiers in selecting ski routes (i.e., individual or combined runs) appropriate to their level of ability, preferences, and current mood and track conditions. The system first asks the user to indicate her ski-run preferences (e.g., route difficulty level), and then uses the indicated preferences to compute a list of recommended routes. The system visualizes on the map the recommended routes and their suitability for the user.

Our approach differs from those used in these mentioned systems in four aspects. First, in our approach, when the interaction starts the user is supported to specify initial preferences, but this initial user input is optional, i.e., the user can obtain some recommendations without explicitly providing initial input. The idea is that the user's preferences may be completely acquired during the product search interaction, using the critiquing function (see details in Section 3). Second, in these mentioned systems, the user is not supported to give feedback to the system's recommendations; or if supported (in [3]), such a feedback is exploited for a future recommendation session, but not for the current one. In our approach, various kinds of feedback are supported: critiquing primarily, but also rating to the selected product. Third, all the mentioned systems do not support users in comparing two recommendations. Fourth, in our approach, case-based reasoning and user-interaction mining methodologies are used to exploit the knowledge contained in past recommendation cases for learning how to rank the recommendations. Among the mentioned systems, only CityGuide [3] learns from past recommendation sessions, exploiting the users' ratings and implicit feedback.

3. The recommendation approach

In this section we discuss our recommendation methodology, identify the limitations of MobyRek's user interface, and present an extended version called MapMobyRek.

In our approach, a recommendation session starts when a user asks for a product suggestion and ends when the user selects a product or when she quits the system. A recommendation session evolves in cycles. In a recommendation cycle, the system shows a ranked list of recommended products (see Fig. 1a) that the user can browse and criticize (see Fig. 1b), and the cycle ends when a new recommendation list is requested and shown.

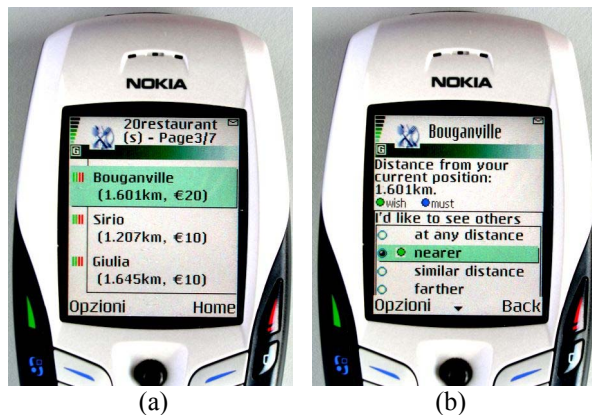


Figure 1. MobyRek user interface
(a) Recommendation list; (b) User critique

Given a user's request asking for a product suggestion, the system integrates the user's initial input, space-time information and long-term preferences to build the user-query representation. The system uses this query representation to compute the recommendation list (see Fig. 1a). If the user is satisfied with a recommended product, then the session ends successfully. However, if the user is somewhat interested in, but not completely satisfied with, a recommended product, then she can criticize that product to express her preferences on the unsatisfactory feature values. When making a critique, the user is supported to indicate the preference strength as a must or a wish condition (see Fig. 1b). This helps the system precisely exploit the user's critique to adapt the query representation. Then, the adapted query representation is used by the system to produce a new recommendation list. When a recommendation session finishes, it is recorded as a case in the system's case base. So, the system can exploit past recommendation sessions in making recommendations for a user in future (more details of our recommendation approach are presented in [12]).

A live-user evaluation of MobyRek showed that it can effectively support mobile users in finding their desired travel products [12]; but this evaluation and some further analysis pointed out several limitations of MobyRek's user interface.

- **Critique influence on the recommendations.** After making a critique, it might be difficult for the user to get an immediate feedback of 1) how the expressed critique influences (changes) the ranking of the recommended items, or 2) the appearance of new recommendations, or 3) the disappearance of previously recommended items. To determine which items lose/gain higher/lower recommendation score, and consequently their rank, in MobyRek the user has to remember the previous result set.
- **Distance perception.** When user is on the move, it is very important to see how far the recommended items are from the user's position. In MobyRek, it is not easy to get an immediate understanding which restaurants are the closest. The user has to scroll down the ranked list of recommended items and compare their distance to her position.
- **Item access inconvenience.** In MobyRek, a user may find it inconvenient when she wants to access an item in a long recommendation list. For instance, if a user is interested in a recommended item and, after browsing other recommendations, decides to come back to see that item, she would have to remember its position in the list or the item name.
- **Items comparison.** A user may be uncertain between two or more recommendations. Hence, a system function providing items comparison could help the user identify the most suitable one(s). This useful function was not supported in MobyRek.

To overcome the above identified limitations and increase the usability of MobyRek, we developed an extended system called MapMobyRek that provides the following new features.

- **Map-based, instead of list-based, visualization.** In MapMobyRek, the recommended items are shown as objects on the map (see Fig. 2a). The map interface supports typical electronic map features, such as zooming in/out or panning up/down/left/right. On the map, a user can easily perceive the recommended items' relative distance to her position and directly access any of the recommended items.
- **Color encoding to represent the recommendation level.** MapMobyRek uses a range of colors to show the predicted degree of the suitability of the recommended items. The colors range is red-orange-yellow-green, where the best recommendations are shown in green and the worst in red. In other words, the ranked list produced by the recommendation algorithm is divided into four parts, where the items in the highest part of the list are encoded as green, those in the second part yellow, and so on. Smiley icons, shown on the top of the mobile screen, help

the user understand the meaning of these colors (see Fig. 2a).

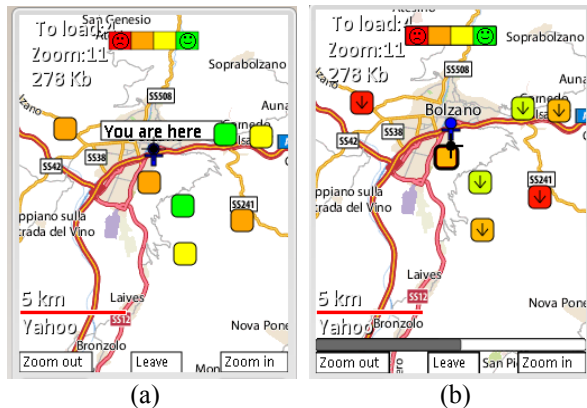


Figure 2. MapMobyRek user interface
(a) Visualization of the recommendations; (b) Recommendation-level change after a critique

- Visualization of the user critique's influence on the recommendations.** In our recommendation methodology, a critique stated as a must condition changes the recommendation list, i.e., some new items are recommended, and some previously recommended items are removed. On the contrary, a critique stated as a wish condition only changes the ranking of the recommended items. In MapMobyRek, the influence of the user critique is visualized immediately and intuitively on the map. After the user has made a wish critique, MapMobyRek changes gradually the colors of the icons on the map (from red to green, or from green to red) and draws an arrow (upward or downward) on each recommended item to show the change (increase or decrease) of its recommendation level (see Fig. 2b). After the user has expressed a must critique, some previously recommended items may not satisfy the new condition, and hence they should be removed from the map. MapMobyRek shows this removal by gradually reducing the icon size of those removed items until they disappear. Besides, MapMobyRek shows the new recommended items by gradually increasing their icon size until the final fixed size.
- Items comparison functionality.** During the interaction with MapMobyRek a user can compare two interested items. After she has selected the two items, their characteristics are displayed side-by-side on the screen; so their advantages and disadvantages can be easily detected (see Figure 3).

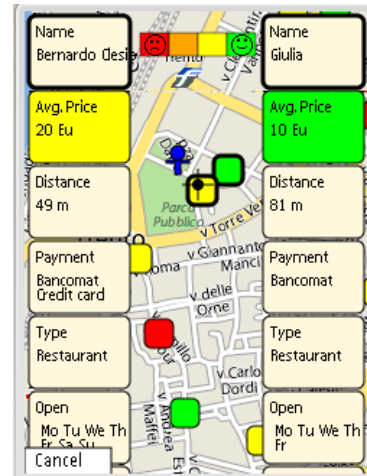


Figure 3. Items comparison functionality

We stress once more time that MobyRek and MapMobyRek implement the same recommendation methodology and ranking algorithm (discussed at the beginning of this section) but they utilize different user interfaces. To support map visualization in MapMobyRek, we used J2meMap (a freeware library for Java MIDP applications) that allows management and display of map-related content. For displaying objects as an overlay of the map, we used the J2ME SVG (Scalable Vector Graphics) library to draw various geometrical shapes and text objects.

4. Research hypotheses and evaluation procedure

We hypothesize that MapMobyRek, compared with MobyRek, provides the following benefits:

- increase the system usability and the user satisfaction,
- reduce the interaction length (i.e., the number of recommendation cycles),
- reduce the time to complete the task (i.e., finding a suitable product), and
- MapMobyRek users tend to choose the items with high recommendation degree, i.e., with a green color. These are the items at high positions (ranks) in the ranked list computed by the recommendation algorithm (see [12] for details in the ranking of the recommended items).

To validate our research hypotheses we ran a user evaluation, which involved twenty testers, using a Nokia N95 mobile phone. In this experiment each tester was asked to use both systems, MobyRek and MapMobyRek, and express his/her subjective opinions. Each tester used both systems (one after the other) to find her desired restaurant and add it to her

travel notes (i.e., a container of the travel products, services, and information selected by the user for her trip). The order in which the two systems were used was assigned randomly to compensate any learning effect, i.e., users may perform the task better with the latter system because they get more familiar with the product domain.

In the comparative evaluation of the two systems, we collected some objective and subjective measures. The objective measures consisted of:

- the average interaction length (i.e., the average number of recommendation cycles – a cycle is defined by the application of a critique),
- the average time to complete the task, and
- the average position of the selected restaurant in the ranked list that was computed and used to generate the color encoding.

To collect a subjective evaluation of the two systems, the tester, after she has used each system, was asked to complete a usability questionnaire for the tested system (see Table 1). In particular, the tester was asked to judge the statements in the questionnaire using a 5-point scale, where 1 means “strongly disagree” and 5 “strongly agree”.

Table 1. The usability questionnaire

ID	Statements
S1	It was simple to use this system.
S2	It was easy to find the information I needed.
S3	The information (such as on-screen messages, and other documentation) provided by this system was clear.
S4	It was easy to learn to use this system.
S5	The information is effective in helping me complete the task and scenario.
S6	The interface of this system is pleasant.
S7	I found it easy to understand what the best recommended items are.
S8	I found it useful to criticize a restaurant and get new offers.
S9	This system has all the functions and capabilities I expected.
S10	Overall, I am satisfied with this system.

Finally, after the tester has completed the systems’ test, she was asked to complete a final questionnaire to express her preference between MobyRek and MapMobyRek (see Table 2). (A similar test approach has been used in the experiment discussed in [10].)

Table 2. The final preference questionnaire

ID	Questions
Q1	What system do you find more informative?
Q2	What system has the better interface?
Q3	What system do you find more useful?
Q4	What system do you prefer?

5. The evaluation results

When mining the log file, which recorded the testers’ recommendation sessions data, we found that in both systems the majority of the testers could find their desired restaurant within 2-3 recommendation cycles. Furthermore, the average interaction length in MapMobyRek was 17% shorter than that in MobyRek (see Fig. 4), which proves our hypothesis that MapMobyRek improves (reduces) the interaction length.




Figure 4. Comparison on the average interaction length

Regarding the average time needed to accomplish the task, MapMobyRek takes 1.43 minute longer than MobyRek (see Fig. 5). This result seems to reject our hypothesis that MapMobyRek users can accomplish the task faster than MobyRek ones. However, we observed that in MapMobyRek the time required to download the maps to the phone caused significant delays. Hence, we suppose that this increment of the time to complete the task was not due to an additional cognitive load asked to the user and could be reduced by a better implementation of the map management functionality (e.g., caching map data and with a better bandwidth).




Figure 5. Comparison on the average time to complete the task (hh:mm:ss)

The average position of the selected item in the final recommendation list is an indication of the goodness of the ranking algorithm. In other words, if the user on average selects items with high recommendation degree, i.e., those mapped to the green color, it means that either the user agrees with the recommendation or she is influenced by the system and follows its suggestion. The average position of the selected restaurant was high (2.65) for MobyRek and even higher (2.3) for MapMobyRek (see Fig. 6). This is quite surprising, given the fact that MapMobyRek does not show precisely the items' positions in the recommendation list, and just uses a rough scale (4 levels, corresponding to the colors). So, we explain this by conjecturing that MapMobyRek users were strongly influenced by the color encoding and that MapMobyRek users selected the "green" restaurants on the maps more often than MobyRek users selected the top restaurants in the ranked lists.




Figure 6. Comparison on the average position of the selected restaurant

Fig. 7 shows the comparison of the testers' average rating to the statements contained in the usability questionnaire. As shown in Fig. 7, both systems received a positive evaluation. However, overall MapMobyRek received higher rating values than MobyRek. We performed a two-tailed, paired t-test on each statement's results looking for significant differences. We found that the significant differences between the two systems relate to the statements S2,

S3, S6, S7, S9, and S10. This means that MapMobyRek provides an easier information search tool, it has a clearer information display and a more pleasant interface, it is more transparent for the users, it has a richer set of system functionalities, and it is generally preferred by the users. Nonetheless, for some statements, such as S1 and S8, the subjective results presented almost no differences between the two systems. For instance, the result for statement S8 ("it is useful to criticize a restaurant and get new offers") shows that the map-based interface does not change the user evaluation on this aspect. This is correct because the same critiquing functionality is implemented in both systems.

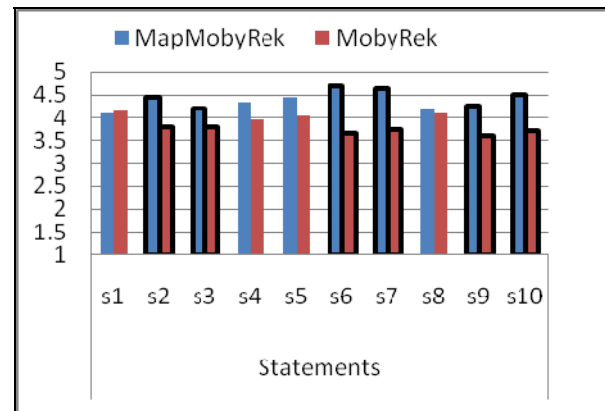


Figure 7. Comparison on the statements in the usability questionnaire

In the final preference questionnaire we asked each tester to vote for the preferred system. The comparative results show that almost all the testers stated that the map-based interface is better than the list-based one in terms of usefulness, interface and informativeness (see Fig. 8).




Figure 8. Comparison on the questions in the final preference questionnaire

6. Conclusions and future work

In this paper we have presented an approach for integrating recommendation and electronic map technologies to build a map-based mobile recommender system that can effectively and intuitively provide personalized recommendations to mobile users. Our real-user study showed that the map-based interface is more effective than the list-based interface that is typically used in recommender systems. We also found that the integration of a map-based interface in a recommender system increases user satisfaction.

There are several open issues that still must be studied. First, we did not investigate how different mappings of rank values to colors could influence the user decision. It would be interesting to analyze to what extent we can push the concept of “green” (i.e., strongly recommended) item, and how an overabundance of “green” items can influence the user decision. Second, in MapMobyRek the user cannot review the recommendation lists produced in previous cycles. It could be helpful if the system records the recommendation states and supports the user (with an “undo” button) to review a previous recommendation state. Another important topic to investigate is how to integrate recommendations on different product types (e.g., a restaurant and an itinerary) and suggest them as a “package”. This would be very important in a real commercial exploitation of the system since typically services that could generate profit are advertised in conjunction with points of interests that the tourist may like to visit.

7. References

- [1] G. Adomavicius and A. Tuzhilin. Toward the next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE Trans. Knowledge and Data Engineering*, 17(6): 734–749, 2005.
- [2] J. Baus, K. Cheverst, and C. Kray. A Survey of Map-based Mobile Guides. In *L. Meng, A. Zipf, and T. Reichenbacher (eds.): Map-based mobile services: Theories, Methods and Implementations*, pp. 193–209. Springer, 2005.
- [3] M.D. Dunlop, A. Morrison, S. McCallum, P. Ptaskinski, C. Risbey, and F. Stewart. Focussed Palmtop Information Access Combining Starfield Displays with Profile-Based Recommendations. In *Proc. of Workshop on Mobile and Ubiquitous Information Access*, pp. 79–89. Springer, 2004.
- [4] M.D. Dunlop, B. Elsey, and M. Montgomery Masters. Dynamic Visualisation of Ski Data: A Context Aware Mobile Piste Map. In *Proc. of Mobile HCI 2007 Conference*, pp. 211–214. ACM, 2007.
- [5] L. Meng, A. Zipf, and S. Winter. *Map-based Mobile Services: Design, Interaction and Usability*. Springer, 2008.
- [6] D. Mohapatra and S.B. Suma. Survey Of Location based Wireless Services. In *Proc. of the 2005 IEEE Int'l Conference on Personal Wireless Communications*, pp. 358–362. IEEE, 2005.
- [7] Q.N. Nguyen and F. Ricci. Replaying Live-User Interactions in the Off-Line Evaluation of Critique-based Mobile Recommendations. In *Proc. of the 2007 ACM Recommender Systems Conference*, pp. 81–88. ACM, 2007.
- [8] Q.N. Nguyen and F. Ricci. Long-Term and Session-Specific User Preferences in a Mobile Recommender System. In *Proc. of the 2008 Int'l Conference on Intelligent User Interfaces*, pp. 381–384. ACM, 2008.
- [9] G. Pospischil, M. Umlauft, and E. Michlmayr. Designing LoL@, a Mobile Tourist Guide for UMTS. In *Proc. of Mobile HCI 2002 Conference*, pp. 140–154. Springer, 2002.
- [10] J. Reilly, J. Zhang, L. McGinty, P. Pu, and B. Smyth. Evaluating Compound Critiquing Recommenders: A Real-User Study. In *Proc. of the 2007 ACM Conference on Electronic Commerce*, pp. 114–123. ACM, 2007.
- [11] P. Resnick and H.R. Varian. Recommender Systems. *Communications of the ACM*, 40(3):56–58, 1997.
- [12] F. Ricci and Q.N. Nguyen. Acquiring and Revising Preferences in a Critique-based Mobile Recommender System. *IEEE Intelligent Systems*, 22(3): 22–29, 2007.
- [13] C. Steinfield. The Development of Location based Services in Mobile Commerce. In B. Preissl, H. Bouwman and C. Steinfield (eds.): *E-Life after the Dot Com Bust*, pp.177-197. Springer, 2004.
- [14] K. ten Hagen, R. Kramer, M. Hermkes, B. Schumann, and P. Mueller. Semantic Matching and Heuristic Search for a Dynamic Tour Guide. In *Proc. of the Int'l Conference on Information and Communication Technologies in Tourism 2005*, pp. 149–159. Springer, 2005.