STS: A Context-Aware Mobile Recommender System for Places of Interest

Matthias Braunhofer, Mehdi Elahi, and Francesco Ricci

Free University of Bozen-Bolzano, Bozen-Bolzano, Italy
{mbraunhofer,mehdi.elahi,fricci}@unibz.it,
http://www.unibz.it

Abstract. In this demo paper we present a novel context-aware mobile recommender system for places of interest (POIs). Unlike existing systems, which learn users' preferences solely from their past ratings, it considers also their personality - using the Five Factor Model. Personality is acquired by asking users to complete a brief and entertaining questionnaire as part of the registration process, and is then exploited in: (1) an active learning module that actively acquires ratings-in-context for POIs that users are likely to have experienced, hence reducing the stress and annoyance to rate (or skip rating) items that the users don't know; and (2) in the recommendation model that builds up on matrix factorization and therefore can be trained even if the users haven't rated any items yet.

1 Introduction

Tourist's decision making is the outcome of a complex decision process that is affected by "internal" (to the tourist) factors, such as personal motivators or past experience, and "external" factors, e.g., advices, information about the products, or the climate of the destination [12]. Context-aware recommender systems can represent and deal with these influencing factors by extending the traditional two-dimensional user/item model that computes recommendations based only on the ratings given by a community of users to a catalogue of items. This is achieved by augmenting the collected ratings with data about the context of an item consumption and rating [1]. For example, the types of place of interest (POI) that users like can differ significantly depending on whether they are visited on a cold or sunny day. If the system stores, together with the rating, the situation in which a POI was experienced, it can then use this information to provide more appropriate recommendations in the various future target contextual situations of the user.

The first challenge for generating context-aware recommendations is how to identify the contextual factors (e.g., weather) that are truly influencing the ratings and hence that are worth considering [3]. Secondly, acquiring a representative set of in-context ratings (i.e., ratings under various contextual conditions) is clearly more difficult than acquiring context-free ratings. Finally, extending traditional recommender systems to really exploit the additional information brought by in-context ratings, i.e., building more accurate recommendations, is the third challenge for context-aware recommender systems.

In this demo paper, we describe an operational context-aware recommender system, called STS (South Tyrol Suggests). STS is an Android-based mobile application that recommends POIs in South Tyrol (Italy) by exploiting various contextual factors (e.g., weather, time of day, day of week, location, mood) and an extended matrix factorization rating prediction model. STS can generate recommendations adapted to the current contextual situation, for example, by recommending indoor POIs (e.g., museums, churches, castles) on bad weather conditions and outdoor POIs (e.g., lakes, mountain hikes, scenic walks) on good weather conditions. The user's preference model is learned using two different sources of knowledge: (1) personality, in terms of the Five Factor Model, that the system acquires with a simple personality questionnaire, and (2) in context ratings that the system actively collects from the user. This allows - and this is the novel aspect of STS - to personalize recommendations and rating requests even for the new users, by leveraging their personality, which is known to be strongly correlated with their tastes and interests [11].

2 Interaction with the System

This section describes a typical interaction with STS and shows some of its functions. Let's assume a tourist or a citizen who is looking for a POI to visit near to Bozen - Bolzano, Italy. The first run of STS opens the registration screen where the user can enter a username, password, birthdate and gender. After registering into the system, the user is asked to fill out the Five-Item Personality Inventory (FIPI) questionnaire [9], in order to allow the system to assess her Big Five personality traits (i.e., openness, conscientiousness, extroversion, agreeableness, neuroticism) (see Figure 1, left). As an alternative to the FIPI questionnaire, other popular personality questionnaires, such as the 120 or 240 item International Personality Item Pool Representation of the NEO PI-R (IPIP-NEO; see [8]), could have been used. These allow a more accurate and reliable personality assessment. However, these questionnaires are time-consuming, taking at least 10-20 minutes to complete, and hence they are ill-suited for mobile interaction models which are usually short in time or on the move.

The entered birthdate, gender and calculated personality are then used by an active learning component [7], which identifies and requests the user to rate a series of POIs whose ratings are estimated to provide the largest improvement of the quality of the subsequent recommendations (see Figure 1, right). We note that this active learning component is able to provide personalized rating requests, without completely relying on explicit feedback (e.g., ratings) or implicit feedback (e.g., item views) which is usually not available for newly registered users.

After that the system is ready for usage, and the user can browse her personalized recommendations through the main application screen (see Figure 2, left). This screen displays a list of POIs that are considered as highly relevant, consid-

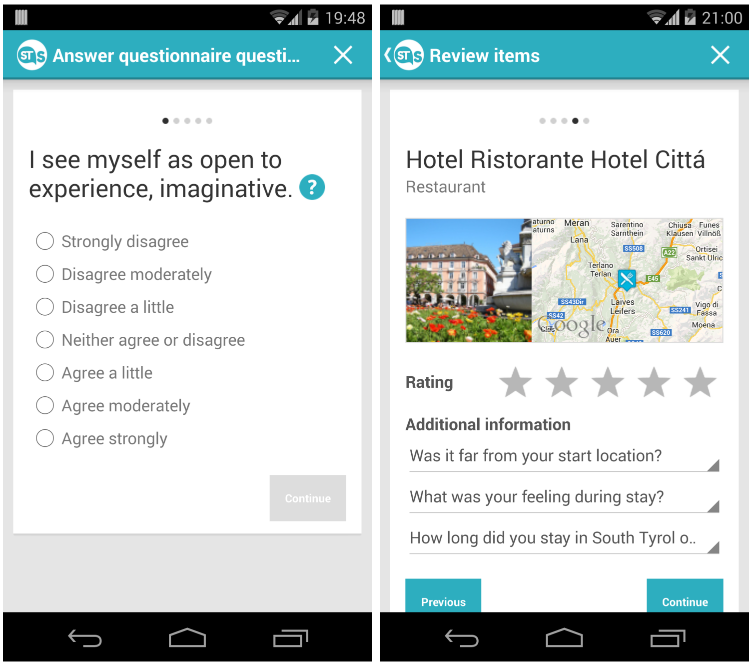

Fig. 1. Screenshots of STS: (left) Five-Item Personality Questionnaire, and, (right) Active Learning

ering the current user's and items' contexts. We note that some context data is automatically acquired by the system (e.g., user's distance to the POIs, weather conditions at the POIs, whereas others can be specified by the user through an appropriate system screen (e.g., user's mood and companion) [5, 4]. In the event the user is interested in one of the POIs, she can click on it and access the POI details window (see Figure 2, right). This window presents various information about the POI, such as a photo, its name, a description, its category as well as an explanation of the recommendation based on the most influential contextual condition. Other supported features include, among others, the ability to write a review for the POI, to view the POI on the map and to bookmark the POI, which then makes it easy to access the POI description.

3 Software Architecture and Implementation

STS implements a rich client always-online architecture, i.e., the client has been kept as thin as possible and it works only in a limited way offline. The client application has been developed using the open-source Android platform, and implements the presentation layer (GUI and a presentation logic). The server application is based on Apache Tomcat server and PostgreSQL database. It

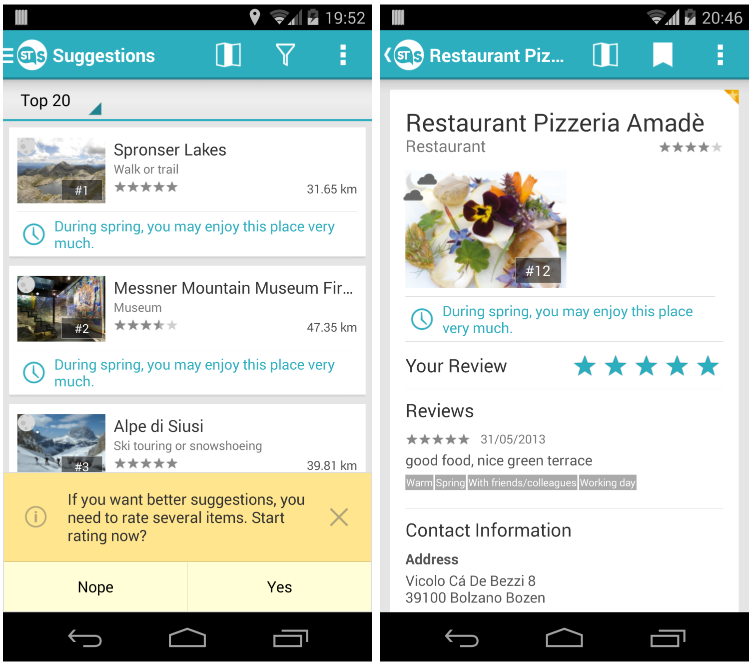

Fig. 2. Screenshots of STS: (left) Recommendation (right) POI description page

implements the data and business logic (recommendation) and makes use of web services and data storages provided by the Regional Association of South Tyrol's Tourism Organizations (LTS[1]), the Municipality of Bolzano[2] and Mondometeo[3] in order to obtain the graphical/textual descriptions as well as weather forecast information for a total of 27,000 POIs. All the server's functionality is exposed via a RESTful web service that accepts and sends JSON objects and that provides several types of resources (e.g., suggestions, POIs, reviews/ratings, user profiles).

4 Recommendation Logic and Evaluation

In order to take into account the current contextual conditions when generating POI recommendations, we have extended the context-aware matrix factorization approach proposed by Baltrunas et al. [3]. This model incorporates baseline parameters for each contextual condition and item (or item category) pair, besides the standard parameters (i.e., global average, item bias, user bias and user-item interaction), in order to capture the deviation of the rating for an item produced

[1] LTS: LTS: http://www.lts.it
[2] Municipality of Bolzano: http://www.comune.bolzano.it
[3] Mondometeo: http://www.mondometeo.org

by the contextual conditions. Since the original context-aware matrix factorization model fails to provide personalized recommendations for users with no or few ratings (i.e., new user problem), we also enhance the representation of a user u by incorporating the set of known user attributes $A(u)$ (i.e., age group, gender and the discretized scores for the Big Five personality traits), analogously as in [10]. A distinct factor vector y_a corresponds to each attribute to describe a user through the set of user-associated attributes $\sum_{a \in A(u)}$. This allows to model the user preferences, even in cases where implicit and explicit feedback are absent.

The resulting model computes a rating prediction for user u and item i in the contextual situation described by the contextual conditions $c_1, ..., c_k$ using the following rule:

$$\hat{r}_{uic_1,...,c_k} = \bar{i} + b_u + \sum_{j=1}^{k} b_{ic_j} + q_i^\top \cdot (p_u + \sum_{a \in A(u)} y_a), \tag{1}$$

where q_i, p_u and y_a are f dimensional real-valued factor vectors representing the item i, the user u and the user attribute a, respectively. $\bar{i}$ is the average rating for item i, b_u is the baseline parameter for user u and b_{ic_j} is the baseline for contextual condition c_j and item i. Model parameters are learned offline, once every five minutes, by minimizing the associated regularized squared error function through stochastic gradient descent.

This recommendation model as well as the implemented active learning strategy for eliciting ratings have been evaluated in two live user studies [5, 7], with the following findings: (1) the recommendation model successfully exploits the weather conditions at POIs and leads to a higher user's perceived recommendation quality and choice satisfaction; and (2) the active learning strategy increases the number of acquired user ratings and the recommendation accuracy in comparison with a state-of-the-art active learning strategy.

5 Conclusions and Future Work

In this demo paper we have illustrated a novel mobile context-aware recommender system, named South Tyrol Suggests (STS), that learns users' preferences from their past ratings as well as their personality. Users's personality is acquired through a brief five-item questionnaire that is subsequently used for actively eliciting ratings for POIs that were estimated to be experienced by users. Finally, this information is exploited for generating high quality recommendations for POIs under the target contextual situation. We have described the implementation of STS, in terms of design, recommendation logic, user interface, and features.

For future work, we plan on making several improvements to STS. Firstly, we intend to provide push recommendations to the user when the current situation seems appropriate, without relying on explicit user's request. Additionally, we would like to exploit in the recommendation process human emotion and the knowledge of the current user activity that can be derived from wearable devices such as smart-watches and smart-bands. Finally, we plan to develop ways to determine the user's personality, without explicitly asking the user to fill a questionnaire by inferring the personality traits of users from their Facebook profiles [2] or their mobile phone usage [6].

References

1. G. Adomavicius, B. Mobasher, F. Ricci, and A. Tuzhilin. Context-aware recommender systems. *AI Magazine*, 32(3):67–80, 2011.
2. Y. Bachrach, M. Kosinski, T. Graepel, P. Kohli, and D. Stillwell. Personality and patterns of facebook usage. In *Proceedings of the 3rd Annual ACM Web Science Conference*, pages 24–32. ACM, 2012.
3. L. Baltrunas, B. Ludwig, S. Peer, and F. Ricci. Context relevance assessment and exploitation in mobile recommender systems. *Personal and Ubiquitous Computing*, 16(5):507–526, 2012.
4. M. Braunhofer, M. Elahi, M. Ge, and F. Ricci. Context dependent preference acquisition with personality-based active learning in mobile recommender systems. In *Human-Computer Interaction (HCII)*. Springer, 2014 (to be appeared).
5. M. Braunhofer, M. Elahi, F. Ricci, and T. Schievenin. Context-aware points of interest suggestion with dynamic weather data management. In *21st Conference on Information and Communication Technologies in Tourism (ENTER)*. Springer, 2014.
6. R. de Oliveira, A. Karatzoglou, P. Concejero Cerezo, A. Armenta Lopez de Vicuña, and N. Oliver. Towards a psychographic user model from mobile phone usage. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*, pages 2191–2196. ACM, 2011.
7. M. Elahi, M. Braunhofer, F. Ricci, and M. Tkalcic. Personality-based active learning for collaborative filtering recommender systems. In *AI*IA 2013: Advances in Artificial Intelligence: XIIIth International Conference of the Italian Association for Artificial Intelligence, Turin, Italy*. Springer, 2013.
8. L. R. Goldberg, J. A. Johnson, H. W. Eber, R. Hogan, M. C. Ashton, C. R. Cloninger, and H. G. Gough. The international personality item pool and the future of public-domain personality measures. *Journal of Research in Personality*, 40(1):84–96, 2006.
9. S. D. Gosling, P. J. Rentfrow, and W. B. Swann Jr. A very brief measure of the big-five personality domains. *Journal of Research in personality*, 37(6):504–528, 2003.
10. Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
11. P. J. Rentfrow and S. D. Gosling. The do re mi's of everyday life: the structure and personality correlates of music preferences. *Journal of personality and social psychology*, 84(6):1236, 2003.
12. J. Swarbrooke and S. Horner. *Consumer behaviour in tourism.* Routledge, 2007.