Voting Operations for a Group Recommender System in a Distributed User Interface Environment

Wolfgang Wörndl
TU München
Boltzmannstr. 3
85748 Garching
Germany
woerndl@in.tum.de

Part Saelim
TU München
Boltzmannstr. 3
85748 Garching
Germany
part.saelim@tum.de

ABSTRACT

This work investigates distributed user interfaces for group recommender systems. In our scenario of a movie recommender, the user interface is distributed on two platforms: personal mobile devices and a public multi-touch tabletop. Our solution proposes voting operations to better support the consensus building among group members. We have implemented a prototype and conducted a preliminary user study with interesting results.

Categories and Subject Descriptors

H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces – *Collaborative computing*

General Terms

Design, Experimentation, Human Factors.

Keywords

user interfaces, group recommender systems, multi-touch tabletop

1. MOTIVATION AND BACKGROUND

Recommender systems recommend items to an active user or user group based on ratings of items or other information about users and items. For group recommender systems, an additional requirement is to facilitate the agreement process among group member to select an item. For example, users may want to go to a cinema together to watch a movie. They could rate some movies on their personal mobile devices and the system could present recommended items on a shared display such as a multi-touch tabletop. The group can then use the tabletop to refine the query and finally select a movie. A separation in the user interface is also called "distributed user interface" (DUI). Thereby, the components of the user interface are distributed across one or more dimension (input, output, platform, space, and time) [1] with multiple distribution strategies.

Recommendation subtasks for group recommender consist of (1) acquiring preference, (2) generating recommendations, (3) presenting recommendations, and (4) helping group members to arrive at an agreement [2]. Related work includes an early approach by McCarthy et al. [3]. The approach aims to support a group of friends planning a skiing vacation. In our work, the goal is to investigate distributed user interfaces and consensus building for group recommender systems in more detail. In addition to

allow users to rate items and remove recommended movies from the recommendation list, our solution proposes voting operations to help users agreeing on an item. We conducted a user study and present the results after explaining the main design features of our application in this poster paper.

2. SOLUTION DESIGN

The main idea is to let users rate their movies on their personal mobile device (1st subtask in the group recommender procedure) and facilitate the result presentation and consensus building on a multi-touch tabletop (3rd and 4th subtask). Users have the option to browse movies individually and rate them on their mobile device (Fig. 1), or bring existing ratings to the group session. Users can move an item on the tabletop to a dedicated area of the shared display, then the item details will appear on the personal mobile device.
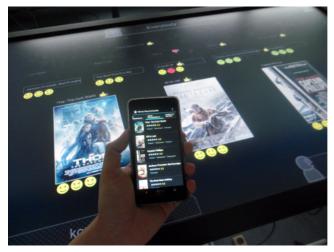

Figure 1. User's view

On the tabletop, the system displays the results of the recommendation process. The three best items are shown with a larger image and additional options are presented with smaller representation (Fig. 1). We refer to the differences in these representations as presentation modes. We use a simple aggregation algorithm based on maximizing average satisfaction as recommendation algorithm. The tabletop provides an overview of recommended items while more detailed information is available on the mobile devices including explanations for recommendations. In a group session, users can *remove* an item from consideration for the current group by moving it to one of the corners of the display. Users can also use a *remove* function on their mobile devices; in this case, an item is removed if more than half of group members submitted the removal.

In addition, we implemented several options to vote on the recommended items. The *promote* operation works like "thumbs-up" buttons commonly found in social networking sites. Similarly, our approach offers a *demote* option. Thus, group members can express how they feel an item is suitable for the current group without having to alter their original personal rating for an item. An important aspect is awareness: users should be able to see the opinions of other group members. This is implemented by smileys (indicating actual or predicted ratings of other users) and thumbs-up/down (for *demote/promote*) icons above and below the movie images in our user interface (Fig. 1). The system also provides a *re-rate* operation on the mobile devices. *re-rate* can be used to change the predicted or actual previous rating of an item, and consequently affect the ranking of movies.

3. USER STUDY

3.1 Setup
We conducted a preliminary user study to evaluate our approach. One of the objectives was to examine how recommendation results can be presented in the multi-device environment in order to help a group of users to arrive at an agreement. We also wanted to investigate which functionalities the group recommender system in this setup should offer on which device.

The first part of the experiment consisted of an explanation and demonstration of all functionalities of the system to the participants. Afterwards, the participants were given the task to determine what movie that want to watch together as a group by using the applications. All user interactions with the system were logged during the sessions. Finally, the participants were asked to complete a survey about the user experience.

To avoid cold start problems or a lengthy initiation procedure, we prepared five stereotype users with initials set of ratings. The users could choose one of the stereotypes and thus start the test session with a reasonable user profile and some movie ratings already present. Since evaluating recommendation accuracy was not one of the goals, this procedure allowed for a quicker setup. The users were of course free to modify existing ratings or add additional ones.

3.2 Results and Conclusion
21 persons participated and were assigned to groups with at least 3 users each. The sample consisted of 11 men and 10 women, ranging between 20 and 30 years old. The participants were mainly students and staff from our university.

Overall, test users found our approach useful. 16 participants (76%) found the prototype to be "very helpful" or "extremely helpful" in supporting a user group to agree on a movie to watch. One more critical remark was that the number of recommended movies should be lower than our 15 proposed items in three presentation modes. But the three different presentation modes were evaluated as useful in principle.

One question was whether the provision of information on items on both platforms (mobile/tabletop) is redundant. Most participants (71%) believe that it was not redundant and they prefer to keep it available on both devices. 19% of participants (4 of 21) would have preferred to access it on the mobile device only and 10% of participants (2 of 21) preferred to access it on the tabletop only. One participant remarked that the redundancy of

information might depend on the number people in a group interacting with the system at the same time.

remove requests submitted via the tabletop are more influential than via the mobile device, because only the former triggers the immediate removal of a movie. 10 out of the 21 participants agreed with this design choice. 29% of participants (6 of 21) felt that the system should treat the requests from both platforms the same way and 24% of participants (5 of 21) reported that the system should treat the requests from mobile devices as more influential than from tabletop. The other operations (*promote*, *demote* or *re-rate)* were available on the personal mobile devices only. So we asked the participants whether these operations should be provided on the tabletop as well. The results were mixed: 62% of participants felt that this alternative could be either "very helpful" or "extremely helpful" for them. 33% of participants felt that it could be either "not at all helpful" or "slightly helpful" for them.

The participants were asked about the degree of usefulness of each voting operation in terms of supporting them in arriving at a final decision. For each operation, the response could be given on a scale from 1 ("not at all useful") to 5 ("extremely useful"). The *re-rate* operation received the highest mean score of 4.24, while *promote* and *demote* operations got mean scores of 3.67 and 3.62 respectively. The lowest mean score belongs to the *remove* operation of either platform, with a mean score of 3.43 for both options on the mobile device and the tabletop. But some participants noted that the *promote* and *demote* operations required lower effort to perform than to *re-rate* an item.

Our findings indicate that the provided preference awareness plays an important role in helping the participants to arrive at an agreement on a movie. Nearly half of the test users felt that the awareness was "remarkably useful" and none of them found it to be "useless". Some participants noted that they would even like to receive notifications of other group members' activities on their respective mobile devices.

The detailed information of movies is exclusively available on the mobile devices. However, to measure its usefulness, participants were asked about how helpful the information was in terms of facilitating the reaching of agreement on a movie. 14 participants (67%) found the information to be either "very helpful" or "extremely helpful". There were 2, 3 and 2 users who felt the information was "not at all helpful", "slightly helpful" and "moderately helpful" respectively.

Interestingly, when the participants were asked to agree on a movie at the end of the test session, all the groups decided on a movie ranked as one of the top three movies in the group recommendation list. Ongoing and future work focuses on the last past of the introduced group recommender process: the support of consensus building with refined alternatives for voting.

4. REFERENCES

[1] Gallud, J.A., Tesoriero, R, and Penichet, V.M.E. 2011. *Distributed User Interfaces*. Springer, Berlin/Heidelberg.

[2] Jameson, A., and Smyth, B. 2007. Recommendation to groups. *The Adaptive Web*, 596-627, Springer, Berlin/Heidelberg.

[3] McCarthy, K. et al. 2006. CATS: A synchronous approach to collaborative group recommendation. Proc. *FLAIRS Conference*, 86-91, AAAI Press