Preference Elicitation for Group Recommender Systems

Thuy Ngoc Nguyen and Francesco Ricci

Faculty of Computer Science
Free University of Bozen-Bolzano, 39100 Bolzano, Italy
{ngoc.nguyen,fricci}@unibz.it

Abstract. In group decision making, users' behaviour are influenced by their long-term and group-induced preferences. However, how to leverage them is challenging due to their dynamic nature, which is also dependent on the specific group settings. In our work, we employ a group recommendation model that utilizes both types of preferences and we analyze alternative ways of combing them, under diverse group settings. Based on a custom-designed simulation process, we examine the effect of these combinations on the model performance. The experimental results demonstrate that a combination scheme weighing more the long-term preferences is well adapted to the scenarios where the group setting has no impact on users' preferences, but when users tend to be cooperative or when their preferences diverge in the context of groups, users seem to benefit more from a recommender that quicker adapts to the group-induced preferences, which reflect their newly emerging interests.

Keywords: Recommender Systems; Group Recommendations; Conversational Systems; Preference Elicitation.

1 Introduction

Today, the challenge for recommender systems (RSs) is expanding from barely suggesting items that match individual's preferences to recommending those satisfying the needs of a group of users [8]. This is derived from real-life cases where people often participate in activities together with others, e.g., having dinners with friends or traveling with family. Several methods for supporting a group of users in making decisions have been incorporated into group recommender systems (GRSs) [5, 6]. However, most of the existing work in GRSs is based on the assumption that by knowing the individual preferences alone, the GRS should still be able to predict the group choice and generate relevant recommendations. For example, research in [4] provides recommendations to a group by fusing three recommendation techniques: demographic, content-based and collaborative filtering but without considering the users' joint interactions with the system. Conversely, our work assumes that the knowledge of individual preferences prior to a group discussion does not suffice, and the system must track the group discussion to effectively support the group decision making process. In fact, a recent observational study on group decision processes has confirmed that group preferences are constructed during the decision-making process and further stressed that research in GRSs should put more focus on the process itself rather than on solving group recommendation problems in a mechanical way [2].

Motivated by these findings, in our previous work we introduced a group recommendation model that exploits both individual long-term and session-based preferences. The long-term interests are acquired in the form of item ratings. While, the session-based preferences, also known as group-induced preferences, are inferred from users' feedback on items during the group discussion. The model was implemented in a GRS that offers a chat environment in which a variety of decision support and recommendation functions are integrated [7]. The usability and the perceived recommendation quality of the system were evaluated through a controlled live user study. Nevertheless, the user study could not fully assess the system performance, which must be examined under a variety of conditions, which users are likely to experience in a group setting.

In a follow-up study, we have hypothesized that the relative importance of each type of preference could vary according to the specific type of the group that needs to be supported. Hence, we have designed a simulation process in order to analyze the proposed model under different preference combination strategies and in three alternative group dynamics settings that match the three kinds of social impact on users' behavior that have been identified in [3]: (a) *independence* - the group has no effect on the user preferences, (b) *conversion* - the group setting nudges group members to be more similar to each other, and (c) *anti-conformity* - the group setting causes group members to react negatively, so they tend to diverge. In these scenarios, three preference combination strategies have been employed: (i) when the importance of the long-term and session-based preferences is equal, (ii) when a much stronger importance is given to the long-term preferences, and (iii) when greater importance is given to the session-based preferences.

Our intuition suggests that the more the users disclose their session-based preferences, the better the group recommendations become. Thus, we have measured and observed how the utility of the top recommendation changes when the amount of elicited preferences in the group discussion grows. We observed that the proposed model can correctly capture the changes in user preferences and, as more feedback is provided in the simulated group session, the utility of the top recommended item converges to that of the assumed group choice. Moreover, the results show fundamental properties of long-term and session-based preference fusion in group recommendations: in the scenarios (a), a GRS requires less preference information derived from the group discussion while in the scenario (b) and (c) it must take into account more the session-based preferences to faster identify the true preferences of the group.

2 Group Recommendation Logic

In our recommendation model, we monitor and utilize the evolving preferences of users in a group decision making process, and we combine them with long-term preferences that are acquired before the group discussion. Combined users' preferences are modeled and continuously updated in the form of utility functions.

The final recommended items are generated by using the group model that aggregates individual preference models; in our case, we use the *Average* aggregation function.

We call $w^{(u)}$ and $w_G^{(u)}$ the utility vector that represents the preferences of user u expressed *before* and *during* a discussion of group G, respectively. The vector $w^{(u)}$ is determined by using a content-based approach [7]. We call $\phi_G^{(u)}$ the set of constraints on the user u utility function derived from the evaluations given by u in the G group session. In order to infer the user utility function from the constraints in $\phi_G^{(u)}$ we use a technique that was introduced in [9], and previously applied only in conversational RSs for individuals [1].

Next, we search for the utility vector $w_G^{(u)}$ that not only satisfies the inferred constraints in $\phi_G^{(u)}$, but also maximizes the cosine similarity of this vector with the vector $w^{(G)}$, the aggregated utility vector of the group. The resulting optimization problem is formulated as follows:

$$w_G^{(u)} = \arg\max \cos(w_G^{(u)}, w^{(G)}) \text{ s.t. } w_G^{(u)} \text{ sat. } \phi_G^{(u)} \tag{1}$$

Finally, each user utility vector $w^{(u)}$ is updated by taking a linear combination of the long term and short term utility vectors, weighted by the parameter $\sigma \in [0, 1]$, which is the "stability" of the long-term preferences:

$$w^{(u)} = \sigma w^{(u)} + (1 - \sigma) w_G^{(u)} \tag{2}$$

3 Experiments and Results

In the three mentioned scenarios, we have generated user groups and simulated users' behaviors by generating items that they could propose to their group, along with their evaluations for the proposed items. The detailed description of the used dataset can be found at [7]. Afterwards, in each scenario, we have investigated the effect of the stability parameter σ (see Eq. 2) on the recommendation quality, i.e., how the utility of the top recommendation changes as a function of the quantity of user feedback acquired by the system. We recall that the parameter σ is used to balance the preference knowledge elicited before and during the group interaction in three scenarios. Particularly, we considered the following cases: $\sigma = 0.1, 0.5$, and 0.9.

Figure 1 shows how, in the three considered scenarios, the true group utility of the top recommendation converges to the best attainable utility which is the utility of the simulated group choice (the item with the largest group utility according to the true, but unknown utilities of the users).

In the *independence* scenario, if $\sigma = 0.9$, the two utilities are exactly the same. With $\sigma = 0.5$, the utility of the top recommended is optimal after each group member has proposed 7 items, while with $\sigma = 0.1$, the true utility of the top recommendation grows much slower. In general, we notice that if the group members incrementally reveal their preferences, the system is eventually able to learn the user needs and adapt to their requirements.

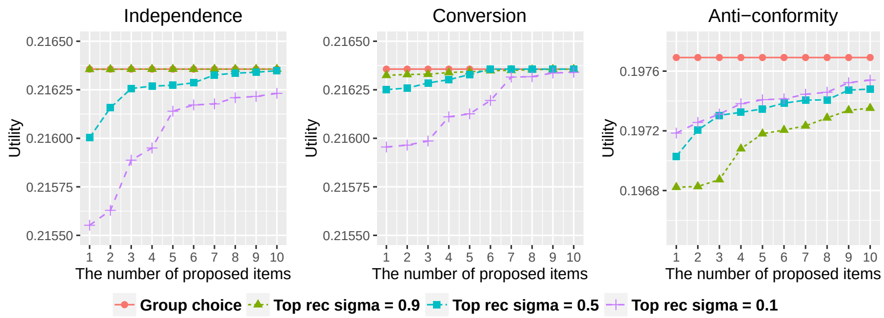

Fig. 1. The true group utility of the top recommended item and the group choice in the three scenarios with respect to random groups of 5 users.

In the *conversion* scenario, when $\sigma = 0.9$, the group utilities of the top recommended item and the group choice are no longer the same as in the *independence* scenario. But, we can see that as the number of proposed items grows, the utility of the top recommended item converges faster to the optimum. This illustrates that if the users have more similar preferences, as expected, the proposed model will learn faster their true preferences.

In the *anti-conformity* scenario, the convergence still occurs but at a slower rate. The results indicate that it is harder to infer the true user profiles if the group members tend to diverge. Additionally, in contrast to the previous scenarios, the group utility of the top recommended item tends to converge more quickly to the optimum when σ is lower.

References

1. Blanco, H., Ricci, F.: Inferring user utility for query revision recommendation. In: Proceedings of the 28th ACM Symposium on Applied Computing. pp. 245–252 (2013)
2. Delic, A., Neidhardt, J., Nguyen, T.N., Ricci, F., Rook, L., Werthner, H., Zanker, M.: Observing group decision making processes. In: Proceedings of the 10th ACM Conference on Recommender Systems. pp. 147–150 (2016)
3. Forsyth, D.R.: Group Dynamics. Wadsworth Cengage Learning, 6th edn. (2014)
4. Garcia, I., Pajares, S., Sebastia, L., Onaindia., E.: Preference elicitation techniques for group recommender systems. Information Sciences 189, 155–175 (2012)
5. Jameson, A., Smyth, B.: Recommendation to groups. The Adaptive Web, LNCS 4321, 596–627 (2007)
6. Masthoff, J.: Group recommender systems: aggregation, satisfaction and group attributes. In: Recommender Systems Handbook. pp. 743–776 (2015)
7. Nguyen, T.N., Ricci, F.: Dynamic elicitation of user preferences in a chat-based group recommender system. In: Proceedings of the 32nd ACM Symposium on Applied Computing. pp. 1685–1692 (2017)
8. Ricci, F., Rokach, L., Shapira, B.: Recommender systems: introduction and challenges. In: Recommender Systems Handbook. pp. 1–34 (2015)
9. Trabelsi, W., Wilson, N., Bridge, D., Ricci, F.: Comparing approaches to preference dominance for conversational recommenders. In: Proceedings of the 22nd IEEE International Conference on Tools with Artificial Intelligence. pp. 113–120 (2010)