Learning the Role and Behavior of Users in Group Decision Making

Dimitris Sacharidis[1], Amra Delic[1], and Julia Neidhardt[1]

E-Commerce Group
Technische Universität Wien, Austria,
{dimitris, neidhardt}@ec.tuwien.ac.at, amra.delic@tuwien.ac.at

Abstract. Group decisions arise in various settings, from mundane everyday tasks such as picking a movie to watch at home, to more involved processes such as that of a hiring committee, and require the interplay among group members, which often leads to compromises. Recommender systems, although primarily designed to cater for individuals, have been extended to support group decision making. However, they are hampered by the lack of models that describe the dynamics in group decisions. In this paper, we present ideas about how to extract more appropriate models for explaining and predicting group decisions, by observing decision outcomes.

Keywords: group recommender systems, group decision making

1 Introduction

People are faced with decisions in their everyday life. The need to make decisions as part of a group may not be that often, but can be considerably more complicated, requiring greater effort and time from all parties involved. Let us consider two separate instances. Assume that a bunch of friends wants to decide on the perfect vacation spot. Each person individually may have different and possibly conflicting preferences and views as to what an ideal holiday is. As a group though, they have to decide on a single location, which unavoidably means that not all preferences or constraints can be accommodated for. In a similar setting, members of a research group decide on a restaurant to have their monthly meeting.

Both examples portray the same set of challenges, namely the task of reconciling individual incompatible preferences. However, they may substantially differ in the way the decision making process is carried out. In the former setting, it is reasonable to assume that all group members play equal roles and that the final decision comes as an "impartial" compromise under mutual understanding. In the latter example, however, this may be far from the truth, as the head of the research group has (usually) stronger influence and authority than the rest.

While recommender systems are typically designed and used to provide recommendations to individual users, there have been ways to extend them in order to offer group recommendations. In literature, there are two main strategies for designing group recommender systems; see also [5] for a comprehensive discussion. The first is to construct a *group profile*, i.e., a collection of item-rating pairs, by merging the profiles of constituent group members. The second is to treat each group member separately, making individual recommendations, and then *fusing recommendations* into a combined list. Notice that, in essence, they both perform an aggregation of ratings, either actual (the profiles in the first strategy) or predicted (the derived recommendations in the second strategy).

We see a missing link here. Existing strategies are well justified given that *no information on the underlying decision making process is available.* Therefore, they have no option but to follow the "least common denominator" approach, e.g., of averaging or aiming for the least misery in the group. On the other hand, understanding and documenting the dynamics of decision making in groups is not an easy task. It is more likely that we can observe and document the outcome of the decision process, rather than its inner workings. Therefore, the question that we raise is whether we can *extract the roles and behaviors that govern group decision making by observing its outcomes.*

In what follows, we first review existing approaches in group recommender systems in Section 2, and outline our approach in Section 3.

2 Past Work

We briefly survey state-of-the-art approaches in group recommender systems.

Group Profile. Each member of the group is associated with a profile; in our setting, a set of item-rating pairs. Methods following this strategy, construct a separate profile for the group containing ratings for every item that a group member has rated. When there are multiple ratings among the group members for a single item, the item's rating in the group is computed using an aggregation function; typically, the average or minimum rating is considered. To provide group recommendations, this merged profile is treated as a virtual user, and then standard user recommendation techniques can be leveraged. A comprehensive survey and evaluation of different strategies can be found in [7].

Recommendation Fusion. Methods under this strategy first employ standard methods to extract recommendation of groups members individually, i.e., no group membership information is used. Then, to predict for a group the rating of an item, they aggregate the individual predicted (or actual) ratings the item is expected to receive (resp. has received) among the group members. Minimum is a popular choice, as it implements the so-called Least Misery strategy. According to it, the predicted group rating of an item should be the least among the member ratings, so as to minimize the chances that any member strongly dislikes the recommendation. Of course, building upon the rich research on rank aggregation methods for Web search [4], other aggregation functions are possible, e.g., Borda count, average, weighted average.

We remark that the Least Misery strategy is rather conservative, since, in the absence of any information about group member roles and behavior, it aims not to displease any member. On the other hand, when member roles/behavior can be assumed to some degree, e.g., all members have equal saying, it makes more sense to average predicted ratings. A more detailed evaluation of techniques aggregating recommendations, with a focus on the goodness of the ranking (instead of on the ratings accuracy), is presented in [2].

Furthermore, regarding the two strategies, we note that the former may be in a disadvantage due to sparse user profiles, which is typically the case. For example, in the extreme setting when no item is commonly rated by any pair of group members, it essentially assigns to the group random profiles with no aggregation. On the other hand, the latter always accounts for all group members, aggregating their predicted ratings. Note though, that in some settings e.g., [3], the former approach was found to be more suitable.

3 Our Approach

Our approach is based on the hypothesis that users in groups play distinct, but unknown, roles that determine the outcome of the decision process. Note that it could be possible that the same person exhibits different behavior in different groups, e.g., in family versus in the work environment. The objective is then to extract the latent roles of group members, in order to explain the observed decisions, and also provide with better group recommendations.

For such an approach to work, we require the availability of decision making outcomes. As a first step, we assume that we have collected the rating of items by groups as a whole. So in essence the data from which we can learn include user-item as well as group-item ratings.

The next step is to model the group decision making process. In the simplest case, we assume that the group rating of an item is determined by a linear combination of individual predicted ratings, where the weights are the parameters modeling user behavior. We then apply standard learning techniques to extract the weights per user and per group. These derived model parameters can then be used to predict the expected rating of items to groups so as to facilitate and expedite decision making.

In the short-term, we plan to perform a user study to help us evaluate the techniques. In particular, we will request people to rate items independently, and also ask them to form groups and provide ratings as a whole. These ratings will be the ground truth upon which we will evaluate how well our techniques can predict group behavior. In addition, we will create synthetic group ratings, where we will simulate different decision making strategies among groups, e.g., average, least misery, maximum pleasure [1].

In the longer term, we plan to learn standard predefined roles for users, e.g., dictator, follower, leader, using classification models that explain the specific behavior of a person in a group. External factors, such as the personality of the user, that persists across groups, may also affect the behavior. Additionally, we plan to consider building richer user behavior models. To achieve this, we will assume an interactive group decision making system and delve into the negotiation process [6]. Particularly, at every information disclosure step, when group members find out the preferences of the others, we would observe the outcome of this step, i.e., how users adapt their preferences to accommodate others.

References

1. Adomavicius, G., Manouselis, N., Kwon, Y.: Multi-criteria recommender systems. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P.B. (eds.) Recommender Systems Handbook, pp. 769–803. Springer (2011), http://dx.doi.org/10.1007/978-0-387-85820-3_24
2. Baltrunas, L., Makcinskas, T., Ricci, F.: Group recommendations with rank aggregation and collaborative filtering. In: RecSys. pp. 119–126 (2010), http://doi.acm.org/10.1145/1864708.1864733
3. Berkovsky, S., Freyne, J.: Group-based recipe recommendations: analysis of data aggregation strategies. In: RecSys. pp. 111–118 (2010), http://doi.acm.org/10.1145/1864708.1864732
4. Dwork, C., Kumar, R., Naor, M., Sivakumar, D.: Rank aggregation methods for the web. In: WWW. pp. 613–622 (2001), http://doi.acm.org/10.1145/371920.372165
5. Jameson, A., Smyth, B.: Recommendation to groups. In: Brusilovsky, P., Kobsa, A., Nejdl, W. (eds.) The Adaptive Web, Methods and Strategies of Web Personalization, pp. 596–627. Springer (2007)
6. Kudenko, D., Bauer, M., Dengler, D.: Group decision making through mediated discussions. In: UM. pp. 238–247 (2003), http://dx.doi.org/10.1007/3-540-44963-9_32
7. Senot, C., Kostadinov, D., Bouzid, M., Picault, J., Aghasaryan, A.: Evaluation of group profiling strategies. In: IJCAI. pp. 2728–2733 (2011), http://ijcai.org/papers11/Papers/IJCAI11-454.pdf