Conversational Group Recommender Systems

Thuy Ngoc Nguyen*
Free University of Bozen-Bolzano
Piazza Domenicani 3
Bolzano, Italy
ngoc.nguyen@unibz.it

ABSTRACT

Recommending to a group of users is multifaceted as people naturally adapt to other members, and it may turn out that what they choose in a group does not fully match individual interests. Besides, it has been shown that the recommendation needs of groups go beyond the aggregation of individual preferences. In practice, it is much more difficult to predict group choices because users take into account the others' reactions and different users react to the group in different ways. Thus, in this research, we aim at exploiting an interactive and conversational approach to facilitate the group decision making process where the complex trade-off between the satisfaction of an individual and the group as a whole typically occurs and needs to be resolved. To attain this goal, we investigate approaches that can access a group situation and autonomously learn an adaptive interaction in a specific condition of the group.

KEYWORDS

Group recommender systems; Group decision processes; Human-computer interaction; User experience; Preference elicitation.

ACM Reference format:
Thuy Ngoc Nguyen. 2017. Conversational Group Recommender Systems. In *Proceedings of UMAP '17, Bratislava, Slovakia, July 09-12, 2017,* 4 pages.
https://doi.org/http://dx.doi.org/10.1145/3079628.3079704

1 INTRODUCTION

Recommender systems (RSs) are tools designed to alleviate information overload by suggesting items that are estimated to fit users' needs and preferences [22]. In many realistic scenarios, the recommended items are consumed by groups of users rather than by individuals [13]. For example, a group of friends or a family may be looking for a restaurant or an attraction site, to go together. The research on group recommender systems (GRSs) is studying methods for supporting a group of users in making decisions when considering a set of alternatives.

Overall, whereas a substantial amount of research in the field of GRSs has focused on group recommendation algorithms, only little research has addressed the role of human-computer interaction in these systems. In particular, most of the previous research has assumed that based on the individual preferences alone, the system can predict a group choice or make good group recommendations. Their primary focus therefore is how to aggregate group members' preferences and identify the "best" items for a group. Conversely, in this research we assume that the knowledge of individual preferences prior to a group discussion does not suffice, and the system must track the group discussion in order to support the group decision making process. This assumption is clearly supported by the fact that there is no clear winner among several preference aggregation techniques that have been proposed in the literature [18], which implies that the group choice depends on the group discussion and not only on the pre-discussion individual preferences. In fact, social scientists studying group dynamics have also stressed the importance of the full decision process adopted by a group in determining the quality of the output decision [10], or in [24], the authors have shown that the degree to which preferences and information are shared within groups, is a key element to understand the group decision making process.

In the context of GRSs, still little attention has been devoted to understanding how the process of making choices in groups can be supported [7]. More concretely, the dynamics of group decision making has been so far under-examined, i.e., users' behavior in the context of a group is overlooked, and the observation of changes in users' preferences during the group decision making process is disregarded. Aside from that, the literature on user experience of RSs has also claimed that although RSs adapt their recommendations to user preferences, they typically do not adapt their interface to support the different decision-making strategies [14]. Driven by these observations, the objective of this research is to support decision making in groups by exploiting the interaction between users and a GRS. This consists of investigating mechanisms that predict a situation where group members are likely to experience in a context of a group, and then provide the most effective supporting actions according to that predicted situation.

2 PROBLEM STATEMENT AND HYPOTHESES

2.1 Situation assessment

Most GRSs developed so far apply a "one-size-fits-all" approach for all group settings, while with each setting, users are likely to react differently. In fact, there are several different kinds of social response to group pressures [10]. For example, group members may be consistent with their personal standards, or show conformity to the group opinion, or alternatively to react negatively to the

*This PhD thesis along with the results presented in this paper are made under the supervision of Francesco Ricci, Free University of Bozen-Bolzano.

UMAP '17, July 09-12, 2017, Bratislava, Slovakia

ACM ISBN 978-1-4503-4635-1/17/07…$15.00
https://doi.org/http://dx.doi.org/10.1145/3079628.3079704

group setting. Motivated by this finding, we conjecture that in different group situations, users' decision-making strategies tend to vary. Thus, it is essential to assess the group situation based on information observed from group settings such as how long a group session lasts and how much users interact with each other. More concretely, we need to deal with the issues: what individual and group features can be detected to assess the group situation and consequently optimize the system support.

2.2 Modeling users' preferences

One of the major challenges for GRSs is to optimally exploit the individual long-term preferences and those induced by the group dynamics called session-based preferences. The session-based preferences that are uncovered in a group discussion could be either consistent or not with the long-term interests that are acquired by the system before the group discussion. Thus, our hypothesis is that the preferences of users should be continuously acquired by observing the evolving behavior of users in the group session, and flexibly integrated with the long-term ones based on a specific situation of the group. Each individual preference model needs to be updated continuously according to the iterative revision of users' preferences during the discussion, and finally aggregated to generate a group preference model.

2.3 Group negotiation support

In most GRSs, users are given group recommendations, and based on these recommendations, they need to negotiate what to do. However, with interactive systems, the challenge is no longer simply recommending items but also guiding and helping users to make informed choices during the negotiation process. We hypothesize that the proactive adaptation of the interaction plays an important role in a group decision support system. In particular, based on the estimated group situation, the system can automatically adapt diverse types of actions, e.g., giving group recommendations, acquiring more information or suggesting a final choice, to support the process of making group decisions.

3 STATE OF THE ART

3.1 Preference aggregation techniques

Two general approaches have been proposed to generate group recommendations: i) *profile aggregation* - aggregate user profiles to create a single profile of the group to which conventional recommendation techniques can be applied, and ii) *recommendation aggregation* - generate individual recommendations for each group member and then combine them to construct a single set of recommendations for the group [13]. The combination of the two approaches, called hybrid switching, has also been exploited in [3]. These approaches have been compared, for instance, in the food recommendation domain [3] or in movie recommendation scenarios [5]. Overall, the choice of which approach is to be used may rely on the domain characteristics, the available data and the precise task. An example of the *profile aggregation* approach is given by *Let's Browse* [15], where an agent assists a group of people in browsing the website by suggesting new material likely to be of common interest. The *recommendation aggregation* method is employed in *Polylens* [21], a system that suggests movies to small groups of people with similar interests. In general, how to optimally aggregate either preferences or ratings or recommendations, is a well-researched topic. In [18], the author gave an overview of different aggregation strategies to reach group decisions. Additionally, the performance of different rank aggregation strategies for generating group recommendations from individual recommendations was investigated by using simulated data of user groups [2].

No matter how the users' preferences are aggregated, most of these approaches have assumed that the personal preferences are adequate to generate group recommendations. In contrast, we believe that the preferences of users are better represented by the combination of both the individual's long-term interests and the session-based preferences.

3.2 Interactive group recommender systems

When it comes to the role of the user-system interaction, research on GRSs has attempted to design interfaces and techniques to support the full decision making process, including the entire preference elicitation and recommendation phases.

The first example is *Intrigue* [1], a tool that assists tour guides in designing tours for heterogeneous tourist groups (e.g., families with children and elderly) by providing recommendations and an interactive agenda. *Travel Decision Forum* [12] allows users to interact with embodied conversational agents representing group members, to define a set of shared preferences, which are discussed and modified by the members themselves. In *Collaborative Advisory Travel System* [19], the concept of critiquing-based RSs is applied, where users can provide feedback in terms of critiques on specific features while the "recommend - review - revise" cycle is repeated until the desired item is found. Also in this direction, *Where2eat* is a mobile app for restaurant recommendation that implements "interactive multi-party critiquing", an extension of the critiquing concept to a computer-mediated conversation between two individuals [11]. *Choicla* [23], a group decision support environment that allows the flexible definition of decision functionality in a domain-free setting. Similarly to critiquing, in *Choicla*, group members are asked to provide evaluations for different item features which are typically specified by a creator, a person who defines a decision task and configures the decision making process. *Hootle+* [17], a GRS that mainly supports the preference elicitation and negotiation process by enabling group members to accept or reject the proposed features and adjust their significance.

While these interactive GRSs mainly support users with recommendations, we speculate that to best support the group decision making process, the system needs to evaluate a group situation based on users' interaction and then automatically and continuously adapt its actions to the estimated condition.

3.3 Evaluation methods

User studies are usually carried out to evaluate the usability of a GRS and the perceived user satisfaction with the recommendations [3, 11, 17, 23]. Offline evaluation studies are limited by the lack of datasets that capture the preferences of users in real group contexts. For this reason, researchers have used data from standard datasets such as MovieLens[1] to test their group recommendation

[1] https://grouplens.org/datasets/movielens/

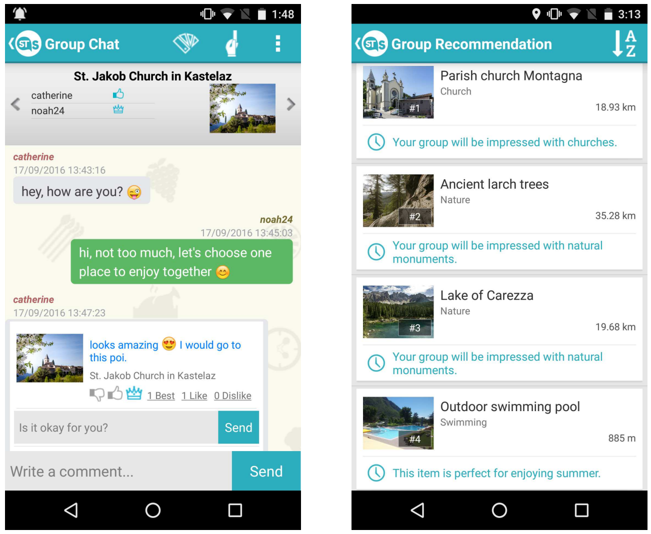

Figure 1: Screenshots of STSGroup, from left to right: (a) Group chat and (b) Group recommendations.

algorithms [2]. Besides, simulation approaches have been used to test conversational RSs [4, 16]. For example, user-system sessions, in which a user incrementally modifies a query to finally select or add a product, were simulated to perform evaluations [16].

In our research, we have developed a chat-based GRS, called STSGroup, which will be introduced later on. With the system, we can collect chat logs that are composed of messages and actions exchanged in real group contexts, and use them to discover possible changes in users' preferences. Moreover, we have designed a procedure that simulates the high degree of interactivity of users with STSGroup under different group situations in order to analyze the performance of the group recommendation model implemented in the system.

4 OVERVIEW OF THE RESEARCH PROGRESS

4.1 Group decision making study

In a first study we explored how people make decisions, and in particular on travel destinations, in a teamwork setting [8, 9]. The study was conducted within a cooperation research project of eleven universities worldwide. Since the study was based on the observations of real user behaviors, it was considered a starting point to understand what group and user features are important for explaining group dynamics and group choices.

The study included three phases: pre-study questionnaire, group meeting/discussion, and post-study questionnaire. The results of this observational study has indicated that group preferences are constructed during the process and stressed that research in GRSs should put more focus on the decision making process taking place in groups rather than on solving group recommendation problems in a mechanical way [9]. This conclusion was supported by the fact that more than two-thirds of the participants, whose the group decision was not in accordance with their most preferred destination, were still satisfied with the collective choice, besides common aggregation strategies in GRSs were hardly able to predict the outcome of the group decision making process.

4.2 Group recommendation model

Interactive GRSs based on critiquing-based techniques suffer from a general drawback: they require users to identify the features that they like or dislike, which can impose a significant cognitive burden on them, especially when the number of features is large [14]. They can also move the discussion on features of the items rather than on the items themselves. Thus, we have proposed a model that acquires users' preferences at the item level, in the form of user's evaluations of the items proposed in a group discussion, and then, based on these evaluations, it automatically determines the item's feature importance weights [20]. The technique, which we adopted to infer the user utility function from constraints on the same utility function by observing what the user likes and dislikes, was introduced in [25], and was applied to conversational RSs for individuals [4]. The individual utility functions consider both individual long-term and session-based preferences, which are aggregated to compute a group utility function that is used to generate group recommendations.

This model was implemented in a GRS to provide a real-time recommendation functionality based on observing a series of user-system interactions. The perceived quality of the group recommendations was evaluated by a user study, and the obtained results showed that the proposed model is able to enhance the perceived group recommendation quality [20].

4.3 Simulating group discussions

While the user study was important to understand whether the proposed recommendation model and GUI are effective and well accepted by users, it was insufficient to thoroughly evaluate its performance, which must be examined under the different conditions that users are likely to experience in a group setting. Thus, we have taken a step further by designing a simulation process to analyze the efficacy of the proposed preference and recommendation model. We studied the effect of alternative settings of the parameter that is used to balance the preference knowledge elicited before and during the group interaction in three user situations. We have considered three kinds of social impacts on users' behavior: (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. In particular, we simulated items that users propose to their group together with their evaluations for other members' proposals in the three situations. It is noteworthy that the proposed items and users' evaluations were simulated differently according to each situation.

We hypothesize that the optimal combination between long-term and session-specific preferences could vary according to specific conditions. We measure the quality of the top-N recommended items for a group, calculate the similarity between the top recommended item and the assumed group choice, and monitor how the utility of the top recommendation changes when the amount of elicited preferences grows. The offline experiments on simulated data have illustrated that the proposed model is able to correctly capture the changes in user preferences and shows some fundamental properties of long-term and session-based preference fusion in group recommendations.

4.4 User interface design

STSGroup (South Tyrol Suggests for Group) is a mobile app that we developed in order to support the process of making decision in groups [20]. It extends the STS app [6], a context-aware places of interest (POIs) recommender originally devoted to individuals. STSGroup targets the discussion stage, where group members' preferences can be elicited and shaped. Particularly, it facilitates the decision making process by allowing group members to join a group chat environment where they can express opinions through text messages and evaluations (see Figure 1a). Based on their actions, the system supports various tasks that the users are likely to undertake during the decision making process, such as comparing options or asking for recommendations (see Figure 1b).

The results of a user study showed that the usability of our system is better than the benchmark and the majority of users found it easy to understand and use.

5 FUTURE WORK

In order to address the first research question (Sec. 2.1), we will investigate techniques for learning individual and group features while observing the users' behavior in various group sessions. Based on the observation and learned characteristics, we will predict the situation of a group. Regarding the second goal (Sec. 2.2), we plan to test and refine the proposed preference revision and aggregation model in group recommendations with more diverse simulation conditions, such as the presence of a dominant member in a group, and with other datasets. For the last research question (Sec. 2.3), we plan to make our current system proactive by adapting the interaction to the specific condition of the group. We also intend to carry out a user study to observe the effects of this adaptivity on user experiences in a group decision making process.

REFERENCES

[1] Liliana Ardissono, Anna Goy, Giovanna Petrone, Marino Segnan, and Pietro Torasso. 2003. Intrigue: personalized recommendation of tourist attractions for desktop and hand held devices. *Applied Artificial Intelligence* 17(8-9) (2003), 687–714.

[2] Linas Baltrunas, Tadas Makcinskas, and Francesco Ricci. 2010. Group recommendations with rank aggregation and collaborative filtering. In *Proceedings of the 4th ACM conference on Recommender systems*. 119–126.

[3] Shlomo Berkovsky and Jill Freyne. 2010. Group-based recipe recommendations: analysis of data aggreagation strategies. In *Proceedings of the 4th ACM conference on Recommender systems*. 111–118.

[4] Henry Blanco and Francesco Ricci. 2013. Inferring user utility for query revision recommendation. In *Proceedings of the 28th ACM Symposium on Applied Computing*. 245–252.

[5] Ludovico Boratto and Salvatore Carta. 2015. The rating prediction task in a group recommender system that automatically detects groups: architectures, algorithms, and performance evaluation. *Journal of Intelligent Information Systems* 45(2) (2015), 221–245.

[6] Matthias Braunhofer, Mehdi Elahi, Francesco Ricci, and Thomas Schievenin. 2013. Context-aware points of interest suggestion with dynamic weather data management. In *Information and communication technologies in tourism 2014*. 87–100.

[7] Li Chen, Marco de Gemmis, Alexander Felfernig, Pasquale Lops, Francesco Ricci, and Giovanni Semeraro. 2013. Human decision making and recommender systems. *ACM Transactions on Interactive Intelligent Systems* 3(3) (2013), 17.

[8] Amra Delic, Julia Neidhardt, Thuy Ngoc Nguyen, and Francesco Ricci. 2016. Research methods for group recommender systems. In *Proceedings of RecTour*.

[9] Amra Delic, Julia Neidhardt, Thuy Ngoc Nguyen, Francesco Ricci, Laurens Rook, Hannes Werthner, and Markus Zanker. 2016. Observing Group Decision Making Processes. In *Proceedings of the 10th ACM Conference on Recommender Systems*. 147–150.

[10] Donelson R. Forsyth. 2014. *Group Dynamics* (6th ed.). Wadsworth Cengage Learning.

[11] Francesca Guzzi, Francesco Ricci, and Robin Burke. 2011. Interactive multi-party critiquing for group recommendation. In *Proceedings of the 5th ACM Conference on Recommender systems*. 265–268.

[12] Anthony Jameson. 2004. More than the sum of its members: challenges for group recommender systems. In *Proceedings of the working conference on Advanced visual interfaces*. 48–54.

[13] Anthony Jameson and Barry Smyth. 2007. Recommendation to Groups. *The Adaptive Web, LNCS* 4321 (2007), 596–627.

[14] Bart P. Knijnenburg, Niels JM Reijmer, and Martijn C. Willemsen. 2011. Each to his own: how different users call for different interaction methods in recommender systems. In *Proceedings of the 5th ACM conference on Recommender systems*. 141–148.

[15] Henry Lieberman, Neil Van Dyke, and Adriana Vivacqua. 1999. Let's browse: a collaborative browsing agent. *Knowledge-Based Systems* 12(8) (1999), 427–431.

[16] Tariq Mahmood and Francesco Ricci. 2007. Learning and adaptivity in interactive recommender systems. In *Proceedings of the 9th international conference on Electronic commerce*. 75–84.

[17] Jesus Omar Alvarez Marquez and Jurgen Ziegler. 2016. Hootle+: A Group Recommender System Supporting Preference Negotiation. In *CYTED-RITOS International Workshop on Groupware*. 151–166.

[18] Judith Masthoff. 2015. Group recommender systems: aggregation, satisfaction and group attributes. In *Recommender Systems Handbook* (2nd ed.), F. Ricci, L. Rokach, and B. Shapira (Eds.). Springer, 743–776.

[19] Kevin McCarthy, Lorraine McGinty, Barry Smyth, and Maria Salamo. 2006. The needs of the many: a case-based group recommender system. In *European Conference on Case-Based Reasoning*. 196–210.

[20] Thuy Ngoc Nguyen and Francesco Ricci. 2017. Dynamic Elicitation of User Preferences in a Chat-Based Group Recommender System. In *Proceedings of the 32nd ACM Symposium on Applied Computing*.

[21] Mark O'connor, Dan Cosley, Joseph A. Konstan, and John Riedl. 2001. PolyLens: a recommender system for groups of users. In *Proceeding of the 7th conference on European Conference on Computer Supported Cooperative Work*. 199 – 218.

[22] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2015. Recommender systems: introduction and challenges. In *Recommender Systems Handbook* (2nd ed.). Springer US, 1–34.

[23] Martin Stettinger, Alexander Felfernig, Gerhard Leitner, Stefan Reiterer, and Michael Jeran. 2015. Counteracting serial position effects in the CHOICLA group decision support environment. In *Proceedings of the 20th International Conference on Intelligent User Interfaces*. 148–157.

[24] R. Scott Tindale and Tatsuya Kameda. 2000. Social sharedness as a unifying theme for information processing in groups. *Group Processes and Intergroup Relations* 3(2) (2000), 123–140.

[25] Walid Trabelsi, Nic Wilson, Derek Bridge, and Francesco Ricci. 2010. Comparing approaches to preference dominance for conversational recommenders. In *Proceedings of the 22nd IEEE International Conference on Tools with Artificial Intelligence*. 113–120.