Double-Sided Recommendations: A Novel Framework for Recommender Systems

Double-Sided Recommendations: a Novel Framework for Recommender Systems

Author1

Department, University
Address
email

Abstract. Recommender systems actively provide users with suggestions of potentially relevant items. In this paper we introduce *double-sided recommendations*, i.e., recommendations consisting of an item and a group of people with whom such an item could be consumed. We identify four specific instances of the double-sided recommendation problem and propose a general method for solving each of them (social comparison-based, group-priority, item-priority and same-priority methods), thus defining a framework for generating double-sided recommendations.
We present the experimental evaluation we carried out, focusing on the restaurant domain as a use case, with the twofold aim of 1) assessing user liking for double-sided recommendations and 2) comparing the four proposed methods, testing our hypothesis that their perceived usefulness varies according to the specific problem instance users are facing. Our results show that users appreciate double-sided recommendations and that all four methods -and, in particular, the group-priority one- can generate useful suggestions.

Key words: double sided recommendations, recommender systems, recommendations to groups, recommendations of groups, social networking.

1 Introduction

In a scenario where the available contents on the Web are constantly growing, *recommender systems* emerge as a specific information filtering technique which actively provides users with suggestions of potentially relevant items, thus helping them to deal with the so called "information overload" problem. Different approaches are usually distinguished based on the information which is needed to generate suggestions: content-based systems employ some knowledge about user preferences and needs, on the one hand, and item features, on the other hand, while collaborative filtering systems base on the opinions of a large community of users.

Most recommender systems target single individuals. In recent years, however, systems which provide recommendations for groups have emerged, based on the idea that some types of recommended items are at least as likely to be used by groups as by individuals, for example vacations, movies, restaurants or cultural events (e.g., concerts or exhibitions).

In many situations, however, a group for which a suggestion could be generated is not necessarily predefined: most of us interact with different individuals, and with different formal and informal groups, from time to time. Also, our social networks usually comprise various (and sometimes overlapping) communities, which relate to different "aspects" of our social lives. For example, we can have friends -and, possibly, different groups of friends-, relatives, school mates, colleagues and sometimes even previous colleagues, as well as occasional acquaintances.

Imagine you are planning to dine out next weekend. Common sense suggests that different restaurants might represent the best choice according to the people you choose to dine. On the other hand, if you are really eager to try a certain (type of) restaurant, some of your contacts may be more willing to accompany you than others. Restaurant and group choices are deeply interconnected. Thus, the most appropriate question you would like a recommender system in the restaurant domain to answer might neither be "Where could I go?", nor "Where could I go with a certain group of friends?", but "Where could I go and with whom?". Similar questions may arise in all other domains where the recommended items are usually enjoyed by groups rather than by single individuals.

To the best of the author's knowledge, no recommender systems exist at present which were designed to answer questions like these. In this paper, we first introduce the idea of recommendations where both an item and a group of people with whom such an item should be consumed are suggested, and call them *double-sided recommendations*. Then, we propose a framework for generating double-sided recommendations.

According to our framework, different instantiations of the double-sided recommendation problem may exist, depending on contextual and occasional elements or on a personal preference for a certain framing of the problem itself. For example, some users might prefer to be recommended an item they can really enjoy, and see the company of other people as an additional treat or just as a way to adhere to unwritten social rules (for example, some individuals might not feel at ease going to a restaurant or to the movies alone; however, they may be quite flexible as far as a company is concerned, provided that it consists of people they like); on the contrary, other users might be primarily interested in spending some time in good company and be ready to compromise on an item which can suit the group as a whole, even if it is not their preferred option.

In our framework, we identify four possible instances of the double-sided recommendation problem, each of which sets different priorities and thus requires a different approach for generating appropriate recommendations. For each problem instance, we provide a generic solution method and some detail on how we exploited it in order to generate double-sided recommendations in our use case. Notice that, in this paper, the focus is on providing a general description of the framework, rather than on discussing specific computational details. Given the basic ideas expressed in the four proposed methods, we believe that different specific techniques might be used in order to compute double-sided recommendations.

An empirical evaluation was carried out in the restaurant domain with the twofold aim of: 1) assessing user liking for double-sided recommendations, demonstrating that these represent a novel and useful service, and 2) comparatively evaluating the four proposed solution methods, as far as their capability of providing *useful* recommendations is concerned.

The paper is structured as follows: Section 2 presents state of the art literature about recommender systems, with a special focus on recommendations to groups, recommendations of groups and social network-based recommendations; Section 3 describes our framework and Section 4 explains how we evaluated our approach and reports our results. Section 5 concludes the paper, with a discussion of possible future work.

2 Related Work

Double-sided recommendations consist of an item and a group of people with whom such an item could be consumed. Thus, they are based on what we could call "single-sided recommendations", i.e., traditional recommendations where only one element -be it an item or a group- is suggested either to a single individual or to a group.

Most literature on recommender systems has focused on the task of recommending potentially relevant items to single individuals. Traditional recommender systems are usually classified according to the information they use in order to assess item relevance and generate recommendations: in *content-based* systems, items are recommended which are similar to those the target user liked in the past [15], in *collaborative filtering* systems, items are recommended which were positively evaluated by users with similar tastes and interests with respect to the target user [17]; finally, in systems which adopt *hybrid approaches*, both types of techniques are used in order to compensate for their respective weaknesses and reach better performances [4].

Some recent approaches have started to consider the social network of the target user as a source of information for generating recommendations, based on the observation that friend-provided suggestions can be more appreciated than those offered by an anonymous system [18], and often proceeding from the consideration that, although many social content sites and recommender systems are appearing which integrate social networking features, no specific guidance is usually provided for selecting interesting items among the huge volume of network-generated contents. Guy et al. [8] found that users prefer recommendations generated taking into account their social network with respect to recommendations based on user-user similarity, as in collaborative filtering, especially when explanations are provided which highlight which people are related to each recommended item. Carmagnola et al. [5] claimed that the mere fact of being part of a social network may cause individuals to modify their attitudes and behaviours because of social influence dynamics, and proposed SoNARS, a recommendation algorithm which explicitly targets users as members of their social network.

Specific issues arise when item recommendations are provided to groups rather than to single individuals. According to Jameson [10], group recommenders are characterized according to 1) the way information about group member preferences is acquired, 2) the way recommendations to groups are generated, 3) the way recommendations are explained (either to individual group members, or to subgroups, or to the group as a whole), and 4) the way group members are eventually helped to achieve consensus. Most related work examines the problem of choosing an appropriate aggregation strategy, depending on the system goals, e.g., maximizing average satifaction or minimizing misery. Recent approaches in recommedations to groups focused on issues related to balancing group and individual satisfaction [11], considering interactions among group members with different personalities (e.g., assertive or cooperative) [16], and explicitly handling disagreement [1].

Finally, a few approaches have considered recommending groups. Most of them focus on suggesting users to affiliate to existing, explicitly-defined communities, based either on structural properties of their social networks (e.g., user proximity to a community [20] or the number and relevance of friends who already belong to it [5]) or on content-related features, such as predicted user interest for the topics which are usually associated to such communities [2]. However, groups to recommend could also be generated on-the-fly by taking into account the social networks of the target users, for example in case a well-matched group of friends to invite to a party should be suggested. Many works in the area of complex network analysis actually focus on the task of identifying relevant subgroups, i.e., sets of nodes (corresponding to individuals in social networks) which are densely connected to each other, while only few links exist which connect them to external nodes. Different approaches exist which either operate in an *a-priori* manner, taking a whole social network as their only input, as in hierarchical clustering-based methods (see for example [14]), or aim at identifying local communities for a given node (see for example [6]). In recent work, methods for finding local communities which contain a set of target nodes [19], and for detecting possibly overlapping communities, thus taking into account the fact that each individual may belong to more than a group [13], are also proposed.

3 Double-Sided Recommendation Framework

We call "double-sided" recommendations where both an item and a group of people with whom such an item should be used are suggested, and formally define them as follows:

DEFINITION (Double-sided recommendation): "Given a target user t, a set of contacts C of the target user and a set of candidate items I, we call a double-sided recommendation either a pair $< i, G >$ where $i \in I$ and $G \subseteq C$; or an N-tuple $< i, G_1, ..., G_{N-1} >$, where $G_n \subseteq C$ and $G_1, ..., G_{N-1}$ are alternative group options, given a certain recommended item; or an N-tuple $< i_1, ..., i_{N-1}, G >$, where $i_n \in I$ and $i_1, ..., i_{N-1}$ are alternative item options, given a certain recommended group."

This definition implies the following assumptions: first, information about user interests or opinions with respect to domain items should be available; otherwise, no item recommendation would be possible. Second, information about the social network of the target user t is needed; otherwise, we would not be able to generate recommendations of groups. More specifically, we assume that a measure of *relationship strength* can be computed in order to assess how important a certain contact c is to the target user t. In the case of our specific implementation, this measure depends on the type and number of actions performed by t which refer to or have an effect on c, such as sending a message or inviting to join a group. A measure of how relevant a group is to the target user is computed as a mean of relationship strength values with respect to all group members.

The proposed definition of double-sided recommendations is very general by design. In fact, it was formulated so that it can encompass all the different situations (i.e., specific *instances* of the double-sided recommendation problem) that users in need of double-sided recommendations may be facing, depending on contextual and occasional elements or on a personal preference for a certain framing of the situation. In our framework, we identify four possible problem instances:

- **Instance 1.** Users are looking for an item to enjoy with some of their contacts. They are very concerned in making a "socially approved" choice and would like to know what the others would do in their place.
- **Instance 2.** Users are interested in spending some time in good company, and they would like to find an item which can please all the people they will meet.
- **Instance 3.** Users are interested in enjoying a pleasant item, and they would like to know who, among their contacts, could keep them company.
- **Instance 4.** Users are interested in enjoying an item in company, and the choice of both a suitable item and a good company are equally important;

Since each specific problem instance sets different priorities, we propose four different methods for solving them: the Social Comparison-based (instance 1), the Group-priority (instance 2), the Item-priority (instance 3), and the Same-priority (instance 4) recommendation method. The last three methods are referred to as *component-based*, since they all base on the identification of structural subcomponents in the social network of the target user (see Section 3.2) in order to generate recommendations of groups.

3.1 Social Comparison-based recommendation method

Method. Taking inspiration from past work on exploiting social influence dynamics in the recommendation process [5] and, in particular, from social comparison theory[1], this method suggests items that were positively evaluated, on average, by relevant others. User relevance depends on both user similarity and user affiliations (i.e., relationship strength), based on the idea that close contacts are more likely to exert some influence. The list of people who evaluated each recommended item is highlighted, in order to leverage social influence. A group recommendation is provided for each item by selecting only the contacts who expressed a positive opinion from such a list.
Detail on how we computed recommendations. In our case, items are recommended based on a threshold value for *itemRelevance*$_i$ = $f(i, R \subseteq C)$, which is a weighted mean of the opinions $itemRelevance_{ic}$ expressed by each contact $c \in R$ of the target user who reviewed item i. In particular, $itemRelevance_{ic}$ = $f^{'}(i, c)$ is computed based on the number, type, and value of actions user c performed on item i. Action *type* is treated as a weight, considering that different types of actions (e.g., rating or bookmarking) may provide different evidence about the strength of user interests for a certain item [12]. As for action *value*, we consider that user actions may have a different polarity (e.g., a

[1] According to social comparison theory, people who are uncertain about what they should be thinking or doing usually seek information about the opinions of relevant others in order to form their own attitudes and behaviours [7].

rating may be positive or negative) and intensity (e.g., a rating of 4 is more positive than a rating of 3). In the case of ratings, the action value corresponds to the rating itself. In the case of actions such as tagging or commenting, action values might be determined by means of some language analysis. However, for simplicity, we decided to determine the value of such actions based on the value of other actions for which it is simply determined (such as rating), assuming that the actions of a certain user on a certain item share the same polarity. A default value is used if no other actions were performed. Finally, the opinion $itemRelevance_{ic}$ of each specific user c is weighted according to the relevance of c to the target user, which is obtained as a mean of his or her scores for relationship strength and similarity with respect to the target user. Similarity depends on user preferences for domain items and is computed based on a variation of the formula for the standard deviation.
Social comparison-based recommendations consisting each of an item, a list of contacts who reviewed it and a recommended group are assigned a recommendation score called *totalScore*, which is the sum of $itemRelevance_i$ and *groupScore*, and are ordered based on it for presentation to the target users.

3.2 Component-based recommendation methods

Methods. The following three methods we propose share three main aspects: 1) individual user preferences for recommendable items are predicted according to a content-based approach; 2) group preferences for items are predicted by aggregating individual group member preferences; 3) recommendable groups are generated based on meaningful substructures which can be identified in the social network of the target user (in our case, connected components) and on simple social rules.
The three proposed methods differ for the facet which is prioritized in providing a double-sided recommendation to the target user: either the item, or the group, or both.

In *group-priority method*, the most relevant groups for the target user are selected at first among all the recommendable ones, based on a threshold score, and then the two best item options are identified for each selected group, according to group preferences.

In *item-priority method*, the best items for the target user are selected at first based on a threshold score for individual preferences, recommendable groups are generated for each item by taking into account only the contacts of the target user whose preference for such an item is higher than a threshold, and then the two best group options are identified according to both group preferences for the item and group relevance to the target user.

In *same-priority method*, all recommendable groups for the target user are combined with all available contents. The best options are selected based on a score which depends on group preferences for the recommended item and group relevance to the target user.
Detail on how we computed recommendations. *Individual preferences* for items are predicted according to a content-based approach. More specifically, a score *itemIndividualScore* = $f^{''}(i,u)$, indicating how interesting item i is expected to be to user u, is computed for every item-user couple, taking into consideration: 1) user u's interests with respect to the domain, 2) overall item interestingness, and 3) specific item interestingness to user u, if available.

User interests are represented in the user model as $<$*feature, value*$>$ pairs, where *feature* corresponds to a category in the reference domain, e.g., "traditional Piedmontese restaurants" for the restaurant domain, and *value* represents the level of interest of user u for that category. In our case, user model values are derived from user actions (such as rating an item) which can be considered indicators for user interest in a certain category. Since each possible item i can be mapped to a category, the level of interest of user u for item i is derived from their interest for the corresponding category.
Overall item interestingness is considered a property of item i, which can be derived from the actions and evaluations of the whole user community. At the moment, only the average user rating is taken into account.
Specific item interestingness to user u is considered a property of the "item-user" pair, which can be derived from the actions and evaluations user u performed on item i (at the moment, we only deal with user bookmarks). We consider specific item interestingness since we assume that favourite items may be included in double-sided recommendations even if users are already aware of them.

Group preferences for items are predicted by aggregating individual preference scores. More specifically, a score *itemGroupScore* = $f^{'''}(i, G)$, indicating how interesting item i is expected to be to group G as a whole, is computed as an average of the individual preference scores of all group members with respect to item i.

Recommendable groups are generated starting from meaningful subsets in the social network of the target user. Such a social network is represented as a graph and may contain either all the contacts of the target user, or, in case an item has already been selected for recommendation, only those contacts whose preference for that item is higher than a threshold. Three sets of connected components, considering respectively family, friendship and all relationships, are identified. Groups are generated from the connected components by: a) eliminating duplicates, and b) applying simple social rules: for example, if the target user has a partner and he or she is not included in a certain recommendable group, another group can be built which also includes him or her. The same rule applies for the target user's best friend, i.e., the contact for whom the value of relationship strength is maximum. Moreover, two groups including, respectively, only the target user's partner and only the target user's best friend are added, if they do not result from connected components.

4 Evaluation

We chose the *restaurant domain* as a use case for our evaluation, considering that restaurants represent a typical example of items which can be recommended to groups, and that people can be assumed to dine with different groups on different occasions.

4.1 Evaluation overview

The evaluation method we adopted consisted of several steps. First, we recruited the experimental subjects among Facebook users, through a snowballing sampling strategy[2].

[2] In snowballing sampling, experimental subjects usually tell the researchers about other individuals who possess the desired characteristics to take part into the study. In our case, it allowed us to recruit people who were connected to each other.

We opted for Facebook users for two reasons: on the one hand, they are accustomed to interactive social websites, and can therefore be considered target users for double-sided recommendations; on the other hand, this allowed us to observe real social networks. 172 people (60% female and 40% male, aged 19-65) accepted to take part to the evaluation at that time. Then, we analyzed their social networks in order to gather information about their social relationships, and their strenght and type (based on Facebook data, we distinguished among “friend”, “family” and “significant other”). This information was stored for successive use.
Experimental subjects were actively involved in the following step, which we call the *opinion gathering phase*: they were asked to use iFOOD[3], an adaptive recommender system in the restaurant domain, for a twenty-day period. This phase aimed at 1) building user models containing information about user interests with respect to different types of restaurants, and at 2) gathering user evaluations of the system contents. 29 subjects out of 172 accepted to take part to this phase (17 female and 12 male, aged 19-62). Information from the opinion gathering phase was combined with information about social networks in order to generate four personalized recommendation lists (one with each of the proposed recommendation methods) for each experimental subject.
In the following step, which we call the *main evaluation phase*, experimental subjects evaluated the double-sided recommendations they were presented, and answered a final short survey. Further detail abouth this phase is provided in the next section.

4.2 Main evaluation phase

This evaluation phase has two goals: assessing user liking for double-sided recommendations and comparatively evaluating the four proposed recommendation methods.
Hypotheses. We hypothesized that users appreciate double-sided recommendations (H1). We also hypothesized that all four methods can provide useful double-sided recommendations (H2); however, we expect that their performances vary according to i) the type of double-sided recommendation problem the experimental subjects are facing during the evaluation (externally-provided problem instance definition, H3), and to ii) the way they usually experience this problem in their real life (personal problem instance definition, H4). In particular, we expect that recommendations generated with a certain method can prove especially useful if users are experiencing the corresponding problem instance (social comparison method is associated to problem instance 1, group-priority method to problem instance 2, item-priority method to problem instance 3 and same-priority method to problem instance 4).
Subjects. All the 29 subjects who took part to the opinion gathering phase were selected as experimental subjects.
Experimental design. Mixed 4 X 5 factorial design, consisting of one within-subject variable (double-sided recommendation method), with four levels (social comparison, group-priority, item-priority and same-priority), and one between-subject variable (externally-provided problem instance definition), with five levels (instance 1, 2, 3, 4, and control situation, where no explicit description of the problem instance subjects are facing is

[3] http://www.piemonte.di.unito.it/progettoDSR/DialogManager?page=home

provided). The experimental subjects were randomly assigned to five groups, corresponding to the five levels of the between-subject variable (9 were assigned to the control group and 5 each to the other four groups).

Material. Recommendations were presented by means of simple web pages. A different page was devoted to each one of the four recommendation lists and navigation was devised so that users could access the following page only after they had completed their tasks on the current one. With the aim of avoiding order-effects, recommendation lists were presented in random order to each experimental subject. Moreover, an initial web page was devoted to the explanation of the experimental task, while a short online survey was presented in the end.

Measures. We evaluated recommendation *usefulness*, i.e., how useful each recommendation is to the experimental subjects in solving the specific problem instance they are facing [9]. We measured recommendation usefulness by means of a 5-point Likert scale, where the first position corresponds to "not useful at all" and the last one to "very useful". Each recommendation was accompanied by the scale to use for its evaluation.
In the survey, *question 1* asked users to express their level of liking for double-sided recommendations in the restaurant domain. *Question 2* aimed at assessing the way users experience the double-sided recommendation problem in their real life (personal problem instance definition): they were provided with four sentences describing the four problem instances we identified in our framework (sentence 2a mapped to instance 1, sentence 2b to instance 2, sentence 2c to instance 3, and sentence 2d to instance 4), and were asked to assess how much each sentence described a situation they experience in their everyday life. All answers should be provided by means of 5-point Likert scales.

Experimental task. Subjects assigned to groups other than the control one were asked to imagine they were facing a specific problem instance, which was described to them according to the instance definitions we provided in our framework presentation (Section 3). All subjects were asked to evaluate at least the first ten recommendations (if available) in each list and to complete the final survey.

Results. Only three subjects out of 29 did not complete the evaluation of all four recommendation lists and did not filled in the final survey. On the whole, 347 evaluations were collected for the same-priority method, 41 for the group-priority method, 284 for the item-priority method and 294 for the social comparison-based method.

User liking for double-sided recommendations. User answers to *question 1* tell us that most subjects were positively impressed by double-sided recommendations (H1): their average rating is 4,38 and 88,4% of subjects expressed a definitely positive opinion, choosing a rating of 4 or 5. A further confirmation of user liking comes from their evaluations of recommendation usefulness: the average evaluation, 3,88, is quite satisfactory. In addition, 67,4% users evaluated the recommendations they received as very useful, rating them 4 or 5.

Comparison among recommendation methods. The average usefulness evaluations are definitely positive for all methods (H2, see Figure 1(b)). However, it seems that the group-priority method is able to provide more useful recommendations, both considering the whole data set and with respect to most problem instances. Thus, we can assume that the group is more relevant than the restaurant when planning to dine out, at least to our experimental subjects. This seems to be confirmed by the fact that subjects tended to

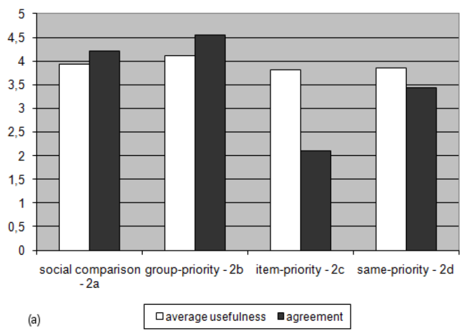

(b)

	Problem Instance					
	1	2	3	4	control	all
Social-comparison	3,71	4,24	3,58	4,27	3,87	3,93
Group-priority	3,83	4,38	4,57	3,6	4,22	4,12
Item-priority	3,31	3,39	4,08	3,8	3,91	3,82
Same-priority	3,23	4,02	4,02	3,95	3,98	3,86

Fig. 1. (a)Comparison between the average usefulness of double-sided recommendation methods and average user agreement with the sentences describing the corresponding problem instances. (b)Average recommendation usefulness evaluations with respect to externally-provided problem instance definitions.

agree the most with sentence 2b (which maps to problem instance 2 and group priority method) in the final survey (see dark bars in Figure 1(a)).

Effect of problem instance. We first consider the *externally-provided problem instance definition* (see Figure 1(b)). Recommendations generated with the group-priority method achieve the highest average usefulness evaluations with respect to three problem instances out of four. However, if we focus on a certain problem instance, it can be observed that the method we expected to have the best performance is always at least second best, as for average recommendation usefulness. Considering the *personal problem instance definition*, we can notice from Figure 1(a) that the values representing the average recommendation usefulness for the four recommendation methods (based on the whole dataset) are quite close to those representing user average agreement with the sentence describing the corresponding problem instance, although user agreement with instance 3 (question 2c) is much lower than the average usefulness of item-priority recommendations. In order to further investigate this issue, we performed a correlational study and found that the average usefulness of recommendations generated with a certain method and for a certain user is positively, although weakly, related to the level of agreement of such a user with the sentence describing the corresponding problem instance (r=0.226, with p = 0,05). Considering both types of problem instance definitions, our results do not yet allow us to validate our hypotheses; however, they seem to support our idea that the specific instance users are facing influences the perceived usefulness of recommendations generated with different methods, even if its effect is not so clearly defined as expected (H3, H4).

5 Conclusion and Future Work

In this paper we have introduced *double-sided recommendations* and proposed a framework, consisting of four different recommendation methods, for generating them.

Results of the experimental evaluation we carried out, focusing on the restaurant domain as a use case, showed that experimental subjects appreciated the possibility of receiving suggestions consisting of a restaurant and a group of people to dine with. Moreover, all the proposed methods proved effective in generating useful recommendations. We also studied the effect of the specific problem instance users are facing on the perceived usefulness of double-sided recommendations generated with different methods, taking into account both externally-provided and personal problem instance definitions. Our results seem to support the idea that some connection between problem instance and perceived recommendation usefulness is actually present, although it is not so clearly defined as it was expected and our data do not yet allow us to statistically confirm our hypothesis.

Although our results refer to a small number of subjects, and should therefore be treated cautiously, we believe that they can be interesting to designers of recommender systems. First, we showed that double-sided recommendations are a novel and useful service: as a consequence, they could be integrated in recommender systems which deal with items which are consumed by groups as often as by individuals (e.g., restaurants, movies or cultural events). As for the specific recommendation methods to use, we found that group-priority method generates particularly useful suggestions, at least in our use case domain. Thus, such a method could be used safely both if it is known that contextual elements and/or personal preferences will determine a specific scenario for group-priority recommendations, and as a default option. The other methods might be used, alone or in conjunction with group-priority method, if it is known that users are facing the specific problem instance they were designed for. However, based on our results, we expect that users will relatively rarely face a situation where they prefer to privilege the item aspect, or where they want to assign the same importance to item and group aspects, in the restaurant domain.

Further research is required in order to investigate the extendability of our results to other domains. For example, it might be possible that methods other than group-priority tend to provide the most useful double-sided recommendations in a different domain. Moreover, an important issue which should be dealt with in future work regards how to determine which type of double-sided recommendation problem users are facing, if the correlation between problem instance and perceived recommendation usefulness is confirmed. For example, this might be inferred from user feedback on recommendations provided with different methods, or from the quality and quantity of their actions on system contents and users (e.g., users who often interact with their social network might be expected to give higher priority to the group aspect).

6 Acknowledgements

This work has been partially supported by Regione Piemonte, grant ICT Converging Technologies 2007, Piemonte Project.

References

1. S. Amer-Yahia, S. Basu Roy, A. Chawla, G. Das, and C. Yu. Group recommendation: Semantics and efficiency. *PVLDB*, 2(1):754–765, 2009.

2. E. Baatarjav, S. Phithakkitnukoon, and R. Dantu. Group recommendation system for facebook. In *OTM '08: Proceedings of the OTM Confederated International Workshops and Posters on On the Move to Meaningful Internet Systems*, pages 211–219, Berlin, Heidelberg, 2008. Springer-Verlag.
3. P. Brusilovsky, A. Kobsa, and W. Nejdl, editors. *The Adaptive Web, Methods and Strategies of Web Personalization*, volume 4321 of *Lecture Notes in Computer Science*. Springer, 2007.
4. R. D. Burke. Hybrid web recommender systems. In Brusilovsky et al. [3], pages 377–408.
5. F. Carmagnola, F. Vernero, and P. Grillo. Sonars: A social networks-based algorithm for social recommender systems. In Geert-Jan Houben, Gord I. McCalla, Fabio Pianesi, and Massimo Zancanaro, editors, *UMAP*, volume 5535 of *Lecture Notes in Computer Science*, pages 223–234. Springer, 2009.
6. J. Chen, O. Zaïane, and R. Goebel. Local community identification in social networks. In *ASONAM '09: Proceedings of the 2009 International Conference on Advances in Social Network Analysis and Mining*, pages 237–242, Washington, DC, USA, 2009. IEEE Computer Society.
7. L. Festinger. A theory of social comparison process. *Human Relations*, 7:117–140, 1954.
8. I. Guy, N. Zwerdling, D. Carmel, I. Ronen, E. Uziel, S. Yogev, and S. Ofek-Koifman. Personalized recommendation of social software items based on social relations. In *RecSys '09: Proceedings of the third ACM conference on Recommender systems*, pages 53–60, New York, NY, USA, 2009. ACM.
9. J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.*, 22:5–53, January 2004.
10. A. Jameson and B. Smyth. Recommendation to groups. In Brusilovsky et al. [3], pages 596–627.
11. J. K. Kim, H. K. Kim, H. Y. Oh, and Y. U. Ryu. A group recommendation system for online communities. *International Journal of Information Management*, 30(3):212 – 219, 2010.
12. A. Kobsa, J. Koenemann, and W. Pohl. Personalised hypermedia presentation techniques for improving online customer relationships. *Knowl. Eng. Rev.*, 16(2):111–155, 2001.
13. A. Lancichinetti, S. Fortunato, and J. Kertsz. Detecting the overlapping and hierarchical community structure in complex networks. *New Journal of Physics*, 11(3):033015, 2009.
14. M. E. Newman. Detecting community structure in networks. *The European Physical Journal B - Condensed Matter and Complex Systems*, 38(2):321–330, March 2004.
15. M. J. Pazzani and D. Billsus. Content-based recommendation systems. In Brusilovsky et al. [3], pages 325–341.
16. J. A. Recio-Garcia, G. Jimenez-Diaz, A. A. Sanchez-Ruiz, and B. Diaz-Agudo. Personality aware recommendations to groups. In *RecSys '09: Proceedings of the third ACM conference on Recommender systems*, pages 325–328, New York, NY, USA, 2009. ACM.
17. J. B. Schafer, D. Frankowski, J. L. Herlocker, and S. Sen. Collaborative filtering recommender systems. In Brusilovsky et al. [3], pages 291–324.
18. R. Sinha and K. Swearingen. Comparing recommendations made by online systems and friends. In *Proceedings of the DELOS-NSF Workshop on Personalization and Recommender Systems in Digital Libraries*, 2001.
19. M. Sozio and A. Gionis. The community-search problem and how to plan a successful cocktail party. In *KDD '10: Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 939–948, New York, NY, USA, 2010. ACM.
20. V. Vasuki, Z. Lu, N. Natarajan, and I. Dhillon. Affiliation recommendation using auxiliary networks. In *RecSys '10: Proceedings of the forth ACM conference on Recommender systems*. ACM, 2010.