Proceedings of the Twenty-Seventh International Florida Artificial Intelligence Research Society Conference

Differential Neighborhood Selection in Memory-Based Group Recommender Systems

Nadia A. Najjar and **David C. Wilson**
University of North Carolina at Charlotte
9201 University City Blvd.
Charlotte, NC 28223
{nanajjar, davils}@uncc.edu

Abstract

As recommender systems have become commonplace to support individual decision making, a need has also been recognized for systems that tailor and provide recommendations to a group of users together rather than individuals alone. Group recommender research to date has focused on evaluating strategies for aggregating profiles of group members to form a consolidated group profile or for aggregating recommendations to individual group members as a consolidated group recommendation list. This paper presents a novel neighborhood selection approach for group recommendation in the context of a neighborhood-based Collaborative Filtering system. We evaluate the performance of this approach with respect to group characteristics such as size and group member similarity. Results show that this approach can result in more accurate predictions for the group, particularly for groups that are more homogenous.

1 Introduction

Recommender systems have traditionally focused on providing decision support that targets an individual end user. As the field has matured, however, the issue of group recommendation has received increasing attention (Jameson and Smyth 2007; Baltrunas, Makcinskas, and Ricci 2010). Group recommender systems balance the preferences of individual members holistically across an entire group of users, in order to create a recommendation that is applicable to the group as a whole. Common group recommendation domains involve social and shared-consumption elements, for example: watching a movie together (O'Connor et al. 2001; Goren-Bar and Glinansky 2004; Senot et al. 2010); eating together (McCarthy 2002; Berkovsky and Freyne 2010; Baltrunas, Makcinskas, and Ricci 2010) or traveling together (McCarthy et al. 2006; Ardissono et al. 2003; Jameson 2004).

Group recommendation has largely been studied in the context of Collaborative Filtering (CF) (Quijano-Sanchez et al. 2013; Berkovsky and Freyne 2010; Baltrunas, Makcinskas, and Ricci 2010), which employs ratings-based user profiles as a foundation for making recommendations. And we focus on group recommendation in the CF context. In traditional individual CF, a user's expected preference for an item is based on the ratings of "neighboring" users for the same item, where neighbors are users that share similar preferences on commonly rated items. Sarwar et al. (Sarwar et al. 2000) divided CF based recommendation into three sub-tasks, *representation*, *neighborhood formation*, and *recommendation generation*.

A considerable amount of research in group-based recommenders has focused on aggregation techniques, and two main group recommendation strategies have been proposed (Jameson and Smyth 2007). The first strategy merges the individual profiles of the group members into one representative group profile (e.g., (Yu et al. 2006; O'Connor et al. 2001)) or *pseudo-user*. The second strategy merges the recommendation lists or predictions computed for each group member into one recommendation list presented to the group (e.g., (Quijano-Sánchez, Recio-García, and Díaz-Agudo 2011; Recio-Garcia et al. 2009)). Particular techniques are commonly inspired by Social Choice Theory, and center around modeling the achievement of consensus among the group (Masthoff 2004). Variations have also been investigated that consider personalities of and social interactions among group members (Gartrell et al. 2010; Recio-Garcia et al. 2009).

Research in profile and predication aggregation strategies has focused on novel aggregation techniques, while employing foundational metrics validated for individual users for underlying neighborhood formation. In our work we explore considering the group as part of the individual prediction generation tasks, in addition to the group aggregation process. We project the tasks defined for a neighborhood-based CF approach into the group-based recommendation process. More specifically we focus on group context as an explicit factor in the *neighborhood formation* and the *recommendation generation* sub-tasks. This paper describes and evaluates our approach to identifying and differentially weighting the contribution of neighbors across the entire group context.

2 Related Work

The major part of research on group recommendation investigated the core group models used for aggregation in generating the group recommendations. Our approach focuses on the recommendation technique itself. In this section we overview some of the work that addresses the group recommendation technique not the aggregation model. Chen et al. (Chen, Cheng, and Chuang 2008) designed a system based on the framework of collaborative filtering. They use a Genetic Algorithm (GA) to exploit known preferences of subgroups of the active group and predict possible similarities among group members. These similarities were used to weight member contributions in item predictions. Their approach to predicting group preferences is based on having access to some item ratings for the target group as well as subgroups of the target group and individual group members' preference information. They use an item-based CF approach to identify items similar to the item under consideration for prediction. If the group did not provide a rating for these items a user-based CF was used to predict the individual ratings. Subgroup information was exploited using using a GA to assign weights in combing the individual users' ratings into a group rating. Then item-based CF was used to calculate the final group rating for the target item.

Berkovsky et al. (Berkovsky and Freyne 2010) investigated the use of aggregated group data in collaborative filtering recipe recommendations. They implemented four weighting models (*uniform, heuristic, role-base, family-log*) for aggregating individual data. The uniform model weights users uniformly, i.e., weight for every user equals 1. The heuristic model is role-based, where a role refers to a user's function within a family: applicant, partner, or child. A user's weight is defined solely by their role. The role-based model weights users according to the activity of users in the same role across the entire community. The family-log model weights users according to their activity in relation to other family members. Extreme case heuristics deal with extremely positive or negative preferences. The least misery heuristic assigns a weight of 1 to the user who provided an extremely negative recipe rating, and a weight of 0 to the other family members. The most pleasure heuristic assigns a weight of 1 to the user who provided the extremely positive data, and 0 otherwise. They evaluated CF recommendations generated using the aggregated data against real-life recipe ratings provided by families interacting with an experimental eHealth portal. The results showed that the most appropriate family-based recipe recommendation strategy should aggregate individual user models, rather than individual recommendations, and weight individual users according to their observed activity rather than according to predefined preferences.

Recio-Garcia et al. (Recio-Garcia et al. 2009) described a group recommender system that takes into account the personality types of the group members. They reported that *Average* and *Least Misery* with personality weighting reflected improvements in the accuracy of the recommendations.

For the neighborhood selection component of a collaborative filtering system we believe accounting for a group neighborhood in calculating the individual predicted ratings will also result in a gain in prediction accuracy. For any group member, when calculating a predicted rating for an item by weighting neighbors ratings, a gain in prediction accuracy is realized by weighting users that belong to the neighborhoods of all group members more than the users belonging only to the individual group member neighborhood.

3 Differential Group Neighborhood Approach

The traditional CF approach is commonly referred to as Neighborhood-based and relies on the fact that each person belongs in a larger group of similarly behaving individuals. For example, items (e.g., products, movies, books, etc.) frequently purchased/liked by the various members of the group can be used to form a basis for recommended items. Similarly, users that appear in more than one of the group members' neighborhoods might be more valuable as a basis for the group recommendation.

Neighborhood-based Collaborative Filtering for single-user recommendation identifies neighbors of the target user and item pair. Extending this to the group-based context, we focus on neighbors of the group as a whole rather than of individual members. In this research we explore the effect on prediction accuracy if special consideration is given to the neighbors of the group members that are shared by all the group members. We hypothesize that prediction accuracy will increase if additional weight is assigned to the neighbors that are common to all group members when used to calculate a predicted rating using the deviation from mean approach as defined in Equation 3 when compared to the baseline predicted rating calculated using Equation 2.

Figure 1 depicts the neighborhoods used in this model. Given an item that the system needs to predict the rating for the group. For each group member we find the *topN* similar users that rated the item. This forms a set of *Neighborhood*s for that item. We then consider the intersection of these neighborhoods. The neighbors that are present across all the individual group members' *Neighborhoods* form the group neighborhood for that item, referred to as the *Group_Neighborhood*. For each group member the users that are present in their *Neighborhood* and not in the *Group_Neighborhood* form what we refer to as the *User_Neighborhood*.

This enables us to distinguish between types of neighbor users in making predictions. In this analysis, we investigate higher weighting for *Group_Neighborhood* users. we use the deviation-from mean-approach to calculate a predicted rating for group members, where the users that are present in the *Group_Neighborhood* are assigned a higher weight than the users that are present in the *User_Neighborhood*. This approach is our *Group Neighborhood Selection* model.

We evaluate the performance of this approach for both main types of group recommendation technique (merging profiles, merging recommendations). We compute predictions for the group based on a pseudo user created by merging the profiles of the group members as well as computing the group predictions based on prediction aggregation. For the prediction aggregation approach a prediction is calculated for each group member and then these individual predictions are aggregated into a final group prediction using a group aggregation model.

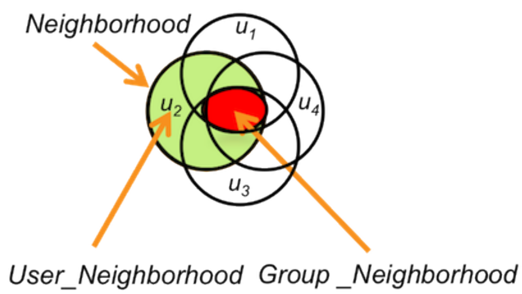

Figure 1: Identifying the neighborhoods

4 Experimental Setup

4.1 Recommendation Technique

For this analysis we implement the most prevalent memory-based CF algorithm, the neighborhood-based CF algorithm (Herlocker, Konstan, and Riedl 2002; Resnick et al. 1994) as a baseline for evaluation. This is employed for the individual user predictions, that are then aggregated for group recommendation as well as the group's pseudo user. The basis for this approach is to calculate a neighborhood similarity between users a and b, w_{ab}, using Pearson correlation:

$$w_{ab} = \frac{\sum_{i=1}^{n}[(r_{ai} - \overline{r}_a)(r_{bi} - \overline{r}_b)]}{\sqrt{\sum_{i=1}^{n}(r_{ai} - \overline{r}_a)^2 \sum_{i=1}^{n}(r_{bi} - \overline{r}_b)^2}} \quad (1)$$

To generate predictions a subset of the nearest neighbors of the active user are chosen based on their correlation. We then calculate a weighted aggregate of their ratings to generate predictions for that user. We use the following formula to calculate the prediction of item i for user a:

$$p_{ai} = \overline{r}_a + \sigma_b \frac{\sum_{b=1}^{n}[(r_{bi} - \overline{r}_b) \cdot w_{ab}]/\sigma_b}{\sum_{b=1}^{n} w_{ab}} \quad (2)$$

Herlocker et al. (Herlocker, Konstan, and Riedl 2002) noted that setting a maximum for the neighborhood size less than 20 negatively affects the accuracy of the recommender systems. They recommend setting a maximum neighborhood size in the range of 20 to 60 users. We set the neighborhood size to 50 users.

For the *Group Neighborhood Selection* model we update the baseline prediction Equation 2 as follows:

$$p_{ai} = \overline{r}_a + \sigma_b \frac{\sum_{b=1}^{n}[(\alpha(r_{bi} - \overline{r}_b) \cdot w_{ab})]/\sigma_b}{\sum_{b=1}^{n} w_{ab}} \quad (3)$$

where α = 1 if neighbor $\in$ *User_Neighborhood*
and α = 2 if neighbor $\in$ *Group_Neighborhood*

4.2 Group Aggregation Strategies

Various group modeling strategies for making recommendations have been proposed and tested to aggregate the individual group user's preferences into a recommendation for the group. Masthoff (Masthoff 2011) evaluated eleven strategies inspired from social choice theory. For clarity, we focus on three representative strategies as a baseline for our evaluation: average strategy, least misery, and most happiness.

Average Strategy This is the basic group aggregation strategy that assumes equal influence among group members and calculates the average rating of the group members for any given item as the predicted rating. Let n be the number of users in a group and r_{ji} be the rating of user *j* for item *i*, then the group rating for item *i* is computed as follows:

$$Gr_i = \frac{\sum_{j=1}^{n} r_{ji}}{n} \quad (4)$$

Least Misery Strategy This aggregation strategy is applicable in situations where the recommender system needs to avoid presenting an item that was really disliked by any of the group members, i.e., that goal is to please the least happy member. The predicted rating is calculated as the lowest rating of for any given item among group members and computed as follows:

$$Gr_i = \min_j r_{ji} \quad (5)$$

Most Happiness This aggregation strategy is the opposite of the least misery strategy. It applies in situations where the group is as happy as their happiest member and computed as follows:

$$Gr_i = \max_j r_{ji} \quad (6)$$

4.3 Pseudo User

To evaluate profile merging group recommendation technique we needed to create a pseudo user for each group by aggregating the profile of the group members. We combine the ratings of the group members, using the average model as defined in Equation 4, where the rating of an item that is rated by one or more group members is added to the pseudo user profile as the average rating for that item over the group members that rated it. For example, if we have a group G_1 of three users $\{u_1, u_2, u_3\}$ with the following (item,rating) profiles: $u_1\{(i_1,5),(i_2,4),(i_3,4),(i_4,1)\}$, $u_2\{(i_1,4),(i_2,3),(i_3,5),(i_5,5)\}$, $u_3\{(i_1,5),(i_3,4),(i_6,4),(i_{10},1)\}$. Then the pseudo group user pG_1 profile would be $\{(i_1,5),(i_2,4),(i_3,4),(i_4,1),(i_5,5),(i_6,4),(i_{10},1)\}$.

4.4 Accuracy Measurement

To evaluate the accuracy of a predicted rating computed for a group across different test conditions, we use the root-mean-square error (RMSE) (Herlocker et al. 2004) RMSE measures the differences between values predicted by a model and the actual values. To do so, we compared the group predicted rating calculated for the test items, using the various implemented approaches to a model of the actual rating (average) across the different group sizes and inner-group similarity levels.

4.5 Data Set and Group Generation

To assess the quality of individual user recommendations, researchers commonly utilize offline evaluations that employ readily available substantial data sets (e.g., Netflix prize[1], MovieLens[2]). This kind of approach can be used

[1]www.netflixprize.com
[2]www.movielens.org

to repeatedly conduct large scale evaluations of proposed techniques. However, when it comes to group-based recommender systems such datasets are not readily available. Generating group-based data directly requires extra overhead in recruiting the groups together and getting them to cooperate and interact towards a common goal at the same time. To address scalability in evaluation, researchers have been utilizing synthetic groups, generated from single-user data sets, to evaluate various approaches to group recommendations (Salamó, McCarthy, and Smyth 2011; Baltrunas, Makcinskas, and Ricci 2010; Amer-yahia et al. 2009; Garcia et al. 2009; Chen, Cheng, and Chuang 2008). The aim here is to develop a standard that has the highest probability of success based on the analysis of the individual characteristics of "real" subjects. Since interactions between individuals are fluid in nature and can always be highly variable having a "real" group versus synthesized can not affect the accuracy of the analysis. We adopt this approach of generating synthetic groups for evaluating our proposed approach to group-based recommendation.

To evaluate the accuracy of an aggregated predicted rating for a group, we use the MovieLens[3] dataset of 1 million ratings from 6040 users on 3882 movies. In creating synthetic groups for evaluation, we varied group size and degree of similarity among group members. The group sizes were varied from 2 to 5. This is a common group size range in this domain. The inner similarity correlation between any two users i, j belonging to group G is calculated using the Pearson Correlation as defined in Equation 1. We defined three inner group similarity levels:

- **high:** 0.5 <= inner group similarity <= 1
- **medium:** 0 <= inner group similarity < 0.5
- **low:** -1 <= inner group similarity < 0

We then randomly created 100 groups for each of the group sizes (2,3,4,5) and similarity levels (high, medium, low), for a total of 1200 unique groups. We placed an additional constraint on group formation, requiring that a valid group must have at least 3 items that were rated by all of the members of the group. This constraint provides for a common group evaluation baseline across those items.

Given the potential for disparity in profile sizes between group members, we employed a training/test set approach to split based on considering individual profile sizes within groups. To generate training and testing sets we split the dataset as follows. For each group we identified the commonly rated items among the group members. Then we checked if that set is larger than 40% of the smallest group member's profile size. If it was smaller then those items would be the testing set for that group. If it was larger then we randomly select items from that set not exceeding 40% of the smallest group member's profile size to compose the testing set for that group. This was to ensure that for each group member we have a majority of their original profile as part of the training set. Once the test items for each group were identified, we created a training and testing set for each group. This ensures that the same training set for each group is used to generate all the predictions for that group. At the end of this we had identified 10,556 group/test item pairs across the 1200 groups.

[3]www.movielens.org

↓ Similarity → Size	**2**	**3**	**4**	**5**
Low	1897	573	390	345
Medium	2158	1117	757	497
High	1368	494	420	540

Table 1: Number of test items across group sizes and similarity levels

To create the training set for each group we start off with the original MovieLens dataset we then add the profile of the pseudo user of that group to the dataset. We then take out the ratings of the test items identified for that group from each of the group members profile and the pseudo user. In other words, the training set for each group is the original MovieLens dataset plus that group's pseudo user profile minus the ratings for the test items for that group for each of the group members and the pseudo user of that group.

We explore outcomes of prediction accuracy for profile merging and recommendation merging using some of the most commonly used group modeling strategies with respect to group size and inner group similarity. For recommendation merging we make a comparison between Average, Least Misery, Most Happiness as defined in Section 4.2. We contrast these approaches using the baseline neighborhood approach to the *Group Neighborhood Selection* model approach. We compare prediction accuracy of the predicted rating of each test item to the average of the actual ratings of the individual group members for that test item.

5 Results and Discussion

We first consider the overall RMSE performance across all group sizes and similarity levels, as shown in Figure 2. We find that the merging profile baseline approach (Pseudo) performs better than any of the baseline aggregating recommendations approaches as well as applying the *Group Neighborhood Selection* model whether we are aggregating profiles or aggregating recommendations.

Given the evaluation approach, it is possible for a subset of the randomly created groups to have a *Group_Neighborhood* size of zero. To examine the potential impact of our model, we specifically consider groups and test items where the *Group_Neighborhood* size is greater than zero. In our results, there were 775 unique groups and 4836 group/item pairs for testing. Figure 3 depicts the results showing that applying the *Group_Neighborhood* model with aggregating profile (Pseudo) provides higher accuracy recommendations ($p<0.05$) than the other implemented approaches.

Group Size To examine how the group neighborhood model performs with respect to the size of the group we combined the evaluated groups based on size. Table 2 shows the RMSE values for the different evaluated models across group sizes 2-5. Results show that for groups of size 2 the average group model (Avg) perfumed best while the group neighborhood model performed better for groups of size 3,

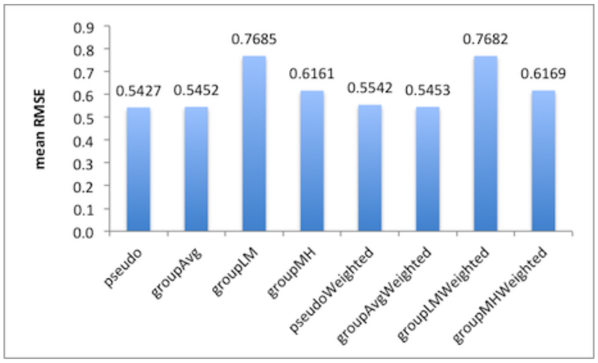

Figure 2: RMSE over all evaluated groups and test items

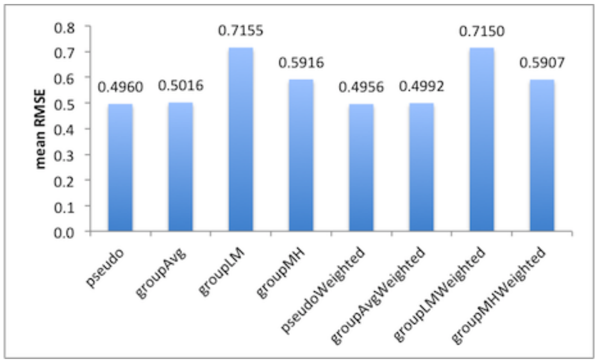

Figure 3: RMSE over all evaluated groups and test items with *Group_Neighborhood* size > 0

4 and 5. For groups of size 3 the average weighted neighborhood model was best while for groups of size 4 and 5 the pseudo weighted neighborhood model had the lowest RMSE.

To directly analyze the performance of the *Group Neighborhood Selection* model we once again examine only the groups/test pairs where it was applicable (*Group_Neighborhood* size >0). These results show that the average prediction aggregation model performed best for groups of size 2 and 3. The average weighted model was best for groups of size 4 and merging profiles approach was best for groups of size 5. From the overall results we conclude that the group neighborhood selection method can achieve significant improvement against other methods in the accuracy of predictions in group-based recommendations.

	pseudo	Avg	LM	MH	pWt	avgWt	lmWt	MhWt
size 2	0.665	**0.653**	0.746	0.688	1.041	0.877	0.805	0.669
size 3	0.570	**0.556**	0.748	0.633	0.571	**0.556**	0.748	0.633
size 4	0.501	0.506	0.757	0.601	**0.500**	0.505	0.757	0.600
size 5	**0.434**	0.465	0.821	0.543	**0.434**	0.465	0.821	0.542

Table 2: RMSE with respect to group size

Group Coherency Examining the performance of the evaluated models with respect to inner group-similarity the group neighborhood method performs best for high similarity groups when used with the aggregating profiles group recommendation technique. This result was not significant

	pseudo	Avg	LM	MH	pWt	avgWt	lmWt	MhWt
High	0.646	0.663	0.818	0.727	**0.644**	0.657	0.817	0.722
Medium	0.531	**0.524**	0.723	0.586	0.533	**0.524**	0.721	0.588
Low	0.495	**0.476**	0.619	0.558	0.500	0.481	0.621	0.566

Table 3: RMSE with respect to inner group coherency for test items with *Group_Neighborhood* size > 0

(p>0.05) when compared to the baseline profile merging approach. For groups with medium and low inner group similarity levels aggregating prediction approach with average group model performs best. We then consider and evaluate only the groups and test items where there was a *Group_Neighborhood* with respect to inner group similarity. Table 3 shows the RMSE values for those groups with test items with *Group_Neighborhood* size > 0. We find that the *Group Neighborhood Selection* model significantly (p<0.05) performs best when used with the aggregating profiles group recommendation technique for groups where the group members are highly similar to each other. For groups with medium to low coherency levels among the group members the average baseline model still performed best.

Group Size and Coherency Another aspect we examined is the combined effect of the group size and the inner group similarity on the prediction accuracy using the proposed approach and the baseline approaches. Table 4 shows the best performing strategy based on the average RMSE calculated for each group size and group inner similarity combination we evaluated.

Similarity↓ Size→	2	3	4	5
Low	Avg	Avg	AvgGNS	Avg
Medium	Avg	AvgGNS	AvgGNS	Pseudo
High	AvgGNS	Pseudo	PseudoGNS	PseudoGNS

Table 4: Best performing technique using RMSE measure

Focusing on the *Group Neighborhood Selection* model, Table 5 shows the best performing approach for these filtered results.

↓ Similarity → Size	2	3	4	5
Low	Avg	Avg	PseudoGNS	Pseudo
Medium	Avg	Avg	Pseudo	Pseudo
High	AvgGNS	PseudoGNS	PseudoGNS	PseudoGNS

Table 5: Best performing technique using RMSE measure for *Group_Neighborhood* size > 0

These results can be summarized that the aggregating predictions using the average group model achieves higher accuracy results for smaller sized groups (2,3) with low to medium similarity. As the group size and inner similarity increases combining the merging profile approach with *Group Neighborhood Selection* model can achieve more accurate results.

6 Conclusion

In various situations the need emerges for recommendations made to a group of people rather than individual users. In this work, we focus of our attention at the selection technique of neighbors in a neighborhood-based CF recombination system tailored to groups. We propose a different selection approach through identifying the neighbors in the selected neighborhood that will be used in the prediction calculation of an item for a group member and happen to appear in all of the other group members' neighborhoods for the same item. We recognize these special neighbors by assigning a higher weight to them than the other neighbors when implementing a deviation-from-mean approach for calculating a prediction for an item. We compared the performance of this approach to some baseline group recommendation models for groups of different group characteristics along the dimensions of size and user-to-user similarity.

Results show that the proposed *Group Neighborhood Selection* model can result in a more accurate prediction when combined with aggregating the profiles of the individual group members to form a *"pseudo"* user for that group. Thus this approach may provide better overall system accuracy, particularly for groups of high inner group similarity. A limitation is that it can only be applied where common neighbors exist between all the group members, though our experimental group formation shows a substantial number of applicable cases. Going further with this work we plan addressing this issue by investigating other models for identifying and weighing group neighbors as part of the prediction computation in group-based recommender systems.

References

Amer-yahia, S.; Roy, S. B.; Chawla, A.; Das, G.; and Yu, C. 2009. Group recommendation: Semantics and efficiency. *Proceedings of The Vldb Endowment* 2.

Ardissono, L.; Goy, A.; Petrone, G.; Segnan, M.; and Torasso, P. 2003. Intrigue: Personalized recommendation of tourist attractions for desktop and hand held devices. *Applied Artificial Intelligence*.

Baltrunas, L.; Makcinskas, T.; and Ricci, F. 2010. Group recommendations with rank aggregation and collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*.

Berkovsky, S., and Freyne, J. 2010. Group-based recipe recommendations: analysis of data aggregation strategies. In *Proceedings of the fourth ACM conference on Recommender systems*, RecSys '10, 111–118. New York, NY, USA: ACM.

Chen, Y.-L.; Cheng, L.-C.; and Chuang, C.-N. 2008. A group recommendation system with consideration of interactions among group members. *Expert Syst. Appl.* 34.

Garcia, I.; Sebastia, L.; Onaindia, E.; and Guzman, C. 2009. A group recommender system for tourist activities. In *Proceedings of the 10th International Conference on E-Commerce and Web Technologies*.

Gartrell, M.; Xing, X.; Lv, Q.; Beach, A.; Han, R.; Mishra, S.; and Seada, K. 2010. Enhancing group recommendation by incorporating social relationship interactions. In *Proceedings of the 16th ACM International Conference on Supporting Group Work*.

Goren-Bar, D., and Glinansky, O. 2004. Fit-recommending tv programs to family members. *Computers & Graphics* 28(2):149 – 156.

Herlocker, J. L.; Konstan, J. A.; Terveen, L. G.; John; and Riedl, T. 2004. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems* 22:5–53.

Herlocker, J.; Konstan, J. A.; and Riedl, J. 2002. An empirical analysis of design choices in neighborhood-based collaborative filtering algorithms. *Inf. Retr.* 5.

Jameson, A., and Smyth, B. 2007. Recommendation to groups. In Brusilovsky, P.; Kobsa, A.; and Nejdl, W., eds., *The adaptive web*.

Jameson, A. 2004. More than the sum of its members: challenges for group recommender systems. In *Proceedings of the working conference on Advanced visual interfaces*.

Masthoff, J. 2004. Group modeling selecting a sequence of television items to suit a group of viewers. *User Modeling and User-Adapted Interaction* 14.

Masthoff, J. 2011. Group recommender systems: Combining individual models. In *Recommender Systems Handbook*.

McCarthy, K.; Salamó, M.; Coyle, L.; McGinty, L.; Smyth, B.; and Nixon, P. 2006. Cats: A synchronous approach to collaborative group recommendation.

McCarthy, J. F. 2002. Pocket restaurantfinder: A situated recommender system for groups. 1–10.

O'Connor, M.; Cosley, D.; Konstan, J. A.; and Riedl, J. 2001. Polylens: a recommender system for groups of users. In *Proceedings of the seventh conference on European Conference on Computer Supported Cooperative Work*.

Quijano-Sanchez, L.; Recio-Garcia, J. A.; Diaz-Agudo, B.; and Jimenez-Diaz, G. 2013. Social factors in group recommender systems. *ACM Trans. Intell. Syst. Technol.* 4(1):8:1–8:30.

Quijano-Sánchez, L.; Recio-García, J. A.; and Díaz-Agudo, B. 2011. Group recommendation methods for social network environments. *3rd Workshop on Recommender Systems and the Social Web* 5th ACM International Conference on Recommender Systems, RecSys'11.

Recio-Garcia, J. A.; Jimenez-Diaz, G.; Sanchez-Ruiz, A. A.; and Diaz-Agudo, B. 2009. Personality aware recommendations to groups. In *Proceedings of the third ACM conference on Recommender systems*.

Resnick, P.; Iacovou, N.; Sushak, M.; Bergstrom, P.; and Riedl, J. 1994. Grouplens: An open architecture for collaborative filtering of netnews. In *1994 ACM Conference on Computer Supported Collaborative Work Conference*.

Salamó, M.; McCarthy, K.; and Smyth, B. 2011. Generating recommendations for consensus negotiation in group personalization services. *Personal and Ubiquitous Computing*.

Sarwar, B.; Karypis, G.; Konstan, J.; and Riedl, J. 2000. Analysis of recommendation algorithms for e-commerce. In *Proceedings of the 2Nd ACM Conference on Electronic Commerce*, EC '00, 158–167. New York, NY, USA: ACM.

Senot, C.; Kostadinov, D.; Bouzid, M.; Picault, J.; Aghasaryan, A.; and Bernier, C. 2010. Analysis of strategies for building group profiles. In *User Modeling, Adaptation, and Personalization*, volume 6075 of *Lecture Notes in Computer Science*.

Yu, Z.; Zhou, X.; Hao, Y.; and Gu, J. 2006. Tv program recommendation for multiple viewers based on user profile merging. *User Modeling and User-Adapted Interaction* 16(1):63–82.