Group Recommender Systems: New Perspectives in the Social Web

Iván Cantador, Pablo Castells

Abstract An increasingly important type of recommender systems comprises those that generate suggestions for groups rather than for individuals. In this chapter, we revise state of the art approaches on group formation, modelling and recommendation, and present challenging problems to be included in the group recommender system research agenda in the context of the Social Web.

1 Introduction

Social Web technologies have emerged as a new step in the course of technological innovations that are having an impact on our everyday lives, reaching the way people relate to each other, work, learn, travel, buy and sell, discover new things, make themselves known, or spend their leisure time. From the common user perspective, while prior technological breakthroughs have to a large extent empowered the individual (giving her instant access to universal online information and services, public authoring access to worldwide publication channels, portable network endpoints, audiovisual production devices, custom-fit adaptation of services to the individual user, etc.), the new trend explicitly emphasises social awareness. As is studied e.g. in Social Sciences, society as a human phenomenon comes along with the notion of group. Groups are indeed a cardinal element in all spheres of social interaction and behaviour. Be they organisations, clubs, political parties, family, tribes, professional units, circles of friends, or just occasional gatherings of people, groups have a part in most human activities, and have played a central role in the evolution of mankind across the ages.

The new social environments open up new possibilities to define, form, articulate, manage and leverage group structures for multiple purposes. The available infrastructure, the explosive growth of online communities, their increasing activity and collected data, lift boundaries and multiply the possibilities to model

I. Cantador (✉), P. Castells
Universidad Autónoma de Madrid 28049 Madrid, Spain
e-mails: ivan.cantador@uam.es, pablo.castells@uam.es

groups as complex units and draw added value from them. Different degrees of group existence can be considered, from sets of people that meet, interact, or have some actual common bond in the physical world, to online contacts that have no relation outside the system, to latent groups of users that are not even directly aware of each other. The new perspectives bring an opportunity to creatively conceive new views and roles for groups in social environments, and perhaps a new angle on the traditional tension between the individual and the group.

As a particular case, in this chapter, we focus on the role of groups in recommender systems. Recommendation technologies are one of the most successful areas of ongoing innovation which find a natural environment to play their best on the Social Web, given the wealth of user input, multiple evidence of user interest, and the huge scale of the new social spaces, where users often count by the million –or billion. Recommender systems have traditionally targeted individual users as the recipients of the personalised system's output. The perspective of delivering shared recommendations for groups as a whole is a new take on the recommendation task that has recently started to be addressed in the field.

The motivation and usefulness of group recommendations naturally arises in situations where a group of users shares a common activity, service, task, or device. For instance, a recommender system could suggest a movie or TV show to watch by a particular group of people (a couple, a family, a group of friends), or could select a sequence of music tracks to be played in a place where individuals with multiple tastes cohabite (a bar, a gym, a shop). Group-oriented recommendation is useful as well in commonplace scenarios such as planning a trip, or choosing a restaurant. Also, in general, many scenarios where ambient intelligence (a.k.a. pervasive/ubiquous computing) technologies take place, and where different people cohabit for a period of time, are susceptible to incorporate group recommendation functionalities. There is a wide number of works in the research literature addressing the group modelling and recommendation problems [16, 19], in different domains and applications, such as recommending tourist attractions [2, 15, 20, 18], food recipes [5], TV programs and movies [4, 10, 23, 27, 36], video clips [18], music tracks and radio stations [8, 21], photos [7], and Web and news pages [25, 29], to name a few.

The group-based perspective changes the recommendation task not only in its purpose, but also in the starting conditions. For instance, the decision of a group member whether or not to accept a given recommendation may depend not only on her own evaluation of the content of the recommendation, but also on her beliefs about the evaluations of the other group members, and about their motivation. As pointed out by Masthoff in [19], the opinion of other members of the group may influence the opinion expressed by a particular user, based on the so-called process of *conformity*, while, on the other hand, the satisfaction of other group members can also lead to increase the user's satisfaction by the so-called *emotional cognition* process. The groups may be quite heterogeneous, in terms of age, gender, intelligence and personality influence on the perception and complacency with the system outputs each member of the groups may have. Thus, a major question that arises is how a recommender system can adapt itself to a group of users, in such a way that each individual enjoys and benefits from the results.

Nowadays, in Web 2.0 systems, people communicate online with contacts through social networks, upload and share multimedia contents, maintain personal bookmarks and blogs, post comments and reviews, rate and tag resources available on the Web, and contribute to wiki-style knowledge bases. The huge amount of user generated content, together with the complexities and dynamics of large groups of people in the Social Web provide room for further research on group recommender systems.

In this chapter, we revise state of the art approaches to group formation, modelling and recommendation, and present challenging problems to be included in the group recommender system research agenda in the context of the Social Web. The rest of the chapter is organised as follows. In Section 2, we explain strategies based on Social Choice Theory, which have been taken into account by the existing group recommendation approaches. In Section 3, we describe group recommender systems presented in the literature, covering different aspects such as group formation, group profile modelling, recommendation aggregation, and cooperative consensus. In Section 4, we revise several open research lines in group recommendation, and in Section 5, we discuss additional challenges that arise in group recommender systems for the Social Web. Finally, in Section 6, we end with some conclusions.

2 Social Choice Theory

Though recommendation approaches have addressed group preference modelling explicitly to a rather limited extent, or in an indirect way in prior work in the computing field, the related issue of *social choice* (also called *group decision making*, i.e. deciding what is best for a group given the opinions of individuals) has been studied extensively in Economics, Politics, Sociology, and Mathematics [24, 33]. The models for the construction of a social welfare function in these works are similar to the group modelling problem we put forward here.

Other areas in which Social Choice Theory has been studied are Collaborative Filtering (CF), Meta-search, and Multi-agent systems. In CF, preferences of a group of individuals are aggregated to produce a predicted preference for somebody outside the group. Meta-search can be seen –and formulated– as a form of group decision making, where the aggregated inputs are produced by information retrieval systems instead of people. In a meta-search engine, the rankings produced by multiple search engines need to be combined into one single list, forming the well-known problem of *rank aggregation* in Information Retrieval [3]. Ensemble recommenders combining several recommendation algorithms also involve a particular case of this problem, similarly to meta-search except for the absence of a query. Finally, in Multi-agent systems, agents need to take decisions that are not only rational from an individual's point of view, but also from a social point of view.

In all the above fields, different strategies to combine several users' preferences and to aggregate item ranking lists can be applied based on the utilised social welfare function. These strategies are classified by Senot and colleagues [27] into three categories, namely *majority-based strategies*, which strength the "most popular" choices (user preferences, item rankings, etc.) among the group, e.g. Borda Count, Copeland Rule, and Plurality Voting strategies; *consensus-based* (or *democratic*) strategies, which average somehow all the available choices, e.g. Additive Utilitarian, Average without Misery, and Fairness strategies; and *borderline* strategies, also called *role-based* strategies in [5], which only consider a subset of choices based on user roles or any other relevant criterion, e.g. Dictatorship, Least Misery and Most Pleasure strategies.

In [17], Mathoff presents and empirically evaluates a number of social choice strategies in a TV item recommendation scenario with a small group of users. Here, we summarise such strategies, and cite representative recommender systems that exploit them. In the following, we assume a user has a preference (utility) for each item represented in the form of a numeric 1-10 rating. In all the cases, the greater the rating value, the most useful the item is for the user.

- **Additive utilitarian strategy**. Preference values from group members are added, and the larger the sum the more influential the item is for the group (Figure 1). Note that the resulting group ranking will be exactly the same as that obtained taking the average of the individual preference values. A potential problem of this strategy is that individuals' opinions tend to be less significant as larger the group is. This strategy could also use a weighted schema, where a weight is attached to individual preferences depending on multiple criteria for single or multiple users. For example, in *INTRIGUE* [2], *weights* are assigned to particular users' ratings depending on the number of people in the group, and the group's members' relevance (children and disabled have a higher relevance).

	Item									
User	i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	10	4	3	6	10	9	6	8	10	8
u_2	1	9	8	9	7	9	6	9	3	8
u_3	10	5	2	7	9	8	5	6	7	6
group	**21**	**18**	**13**	**22**	**26**	**26**	**17**	**23**	**20**	**22**

Fig. 1. Group choice selection following the additive utilitarian strategy. The ranked list of items for the group would be (i_5-i_6, i_8, i_4-i_{10}, i_1, i_9, i_2, i_7, i_3).

- **Multiplicative utilitarian strategy**. Instead of adding the preferences, they are multiplied, and the larger the product the more influential the item is for the group (Figure 2). This strategy could be self-defeating: in a small group, the opinion of each individual may have too much impact on the product.

	Item									
User	***i_1***	***i_2***	***i_3***	***i_4***	***i_5***	***i_6***	***i_7***	***i_8***	***i_9***	***i_{10}***
u_1	10	4	3	6	10	9	6	8	10	8
u_2	1	9	8	9	7	9	6	9	3	8
u_3	10	5	2	7	9	8	5	6	7	6
group	**100**	**180**	**48**	**378**	**630**	**648**	**180**	**432**	**210**	**284**

Fig. 2. Group choice selection following the multiplicative utilitarian strategy. The ranked list of items for the group would be (i_6, i_5, i_8, i_{10}, i_4, i_9, i_2-i_8, i_1, i_3).

- **Average strategy**. In this strategy, the group rating for a particular item is computed as the average rating over all individuals (Figure 3). Note that if no user or item weighting is conducted, the ranking list of this strategy is the same as that of the Utilitarian strategy. *Travel Decision Forum* [15] implements multiple group modelling strategies, including the average strategy and the median strategy, which uses the middle value of the group members' ratings, instead of the average value. In [36], Yu and colleagues present a TV program recommender that performs a variation of the average strategy, where the group preference vector minimises its distance compared to the individual members' preference vectors.

	Item									
User	***i_1***	***i_2***	***i_3***	***i_4***	***i_5***	***i_6***	***i_7***	***i_8***	***i_9***	***i_{10}***
u_1	10	4	3	6	10	9	6	8	10	8
u_2	1	9	8	9	7	9	6	9	3	8
u_3	10	5	2	7	9	8	5	6	7	6
group	**7**	**6**	**4.3**	**7.3**	**8.7**	**8.7**	**5.7**	**7.7**	**6.7**	**7.3**

Fig. 3. Group choice selection following the average strategy. The ranked list of items for the group would be (i_5-i_6, i_8, i_4-i_{10}, i_1, i_9, i_2, i_7, i_3).

- **Average without misery strategy**. As the average strategy, this one assigns an item the average of its ratings in the individual profiles. The difference here is that those items which have a rating under a certain threshold will not be considered in the group recommendations. Figure 4 shows an example of group formation following this strategy with a threshold value of 3. *MusicFX* [21], which chooses a radio station for background music in a fitness centre, follows an average without misery strategy, and a weighted random selection is made from the top stations in order to avoid starvation and always picking the same station. *CATS* system [22] helps users to choose a joint holiday based on individuals' critiques on holiday package features, and applying the misery aspect.

User	Item i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	10	4	3	6	10	9	6	8	10	8
u_2	1	9	8	9	7	9	6	9	3	8
u_3	10	5	2	7	9	8	5	6	7	6
group	**-**	**18**	**-**	**22**	**26**	**26**	**17**	**23**	**-**	**22**

Fig. 4. Group choice selection following the average without misery strategy. The ranked list of items for the group would be (i_5-i_6, i_8, i_4-i_{10}, i_2, i_7).

- **Least misery strategy**. The score of an item in the group profile is the minimum of its ratings in the user profiles. The lower rating the less influential the item is for the group. Thus, a group is as satisfied as its least satisfied member (Figure 5). *PolyLens* [23] uses this strategy, assuming a group of people going to watch a movie together tends to be small, and the group is as happy as its least happy member. Note that a minority of the group could dictate the opinion of the group: although many members like a certain item, if one member really hates it, the preferences associated to it will not appear in the group profile.

User	Item i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	10	4	3	6	10	9	6	8	10	8
u_2	1	9	8	9	7	9	6	9	3	8
u_3	10	5	2	7	9	8	5	6	7	6
group	**1**	**4**	**2**	**6**	**7**	**8**	**5**	**6**	**3**	**6**

Fig. 5. Group choice selection following the least misery strategy. The ranked list of items for the group would be (i_6, i_5, i_4-i_8-i_{10}, i_7, i_2, i_9, i_3, i_1).

- **Most pleasure strategy**. It works as the least misery strategy, but instead of considering for an item the smallest ratings of the users, it selects the greatest ones. The higher rating the more influential the item is for the group, as shown in Figure 6.

User	Item i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	10	4	3	6	10	9	6	8	10	8
u_2	1	9	8	9	7	9	6	9	3	8
u_3	10	5	2	7	9	8	5	6	7	6
group	**10**	**9**	**8**	**9**	**10**	**9**	**6**	**9**	**10**	**8**

Fig. 6. Group choice selection following the least misery strategy. The ranked list of items for the group would be (i_1-i_5-i_9, i_2-i_4-i_6-i_8, i_3-i_{10}, i_7).

- **Fairness strategy**. In this strategy, the items that were rated highest and cause less misery to all the users of the group are combined as follows. A user is randomly selected. His *L* top rated items are taking into account. From them, the item that less misery causes to the group (that from the worst alternatives that has the highest rating) is chosen for the group profile with a score equal to *N*, i.e., the number of items. The process continues in the same way considering the remaining *N*–1, *N*–2, etc. items and uniformly diminishing to 1 the further assigned scores. In the final list, the higher score the more influential the item is for the group. Note that this list would be different if we let other users to choose first. To better understand the strategy, let us explain its first step on the example shown in Figure 7. Suppose we start with user u_1, whose top ranked items are i_1, i_5 and i_9. From these items, we choose item i_5 because it is the one that less misery causes to users u_2 and y_3, whose lowest ratings for items i_1, i_5 and i_9 are respectively 1, 7 and 3. We assign item d_5 a score of 10.

	Item									
User	i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	10	4	3	6	10	9	6	8	10	8
u_2	1	9	8	9	7	9	6	9	3	8
u_3	10	5	2	7	9	8	5	6	7	6
group	**4**	**3**	**1**	**8**	**10**	**9**	**5**	**7**	**2**	**6**

Fig. 7. Group choice selection following the fairness strategy. The ranked list of items for the group could be (i_5, i_6, i_4, i_8, i_{10}, i_7, i_1, i_2, i_9, i_3), following the user selecting order u_1, u_2 and u_3, and setting *L*=3.

- **Plurality voting strategy**. This method follows the same idea of the fairness strategy, but instead of selecting from the *L* top preferences the one that least misery causes to the group, it chooses the alternative which most votes have obtained. Figure 8 shows an example of the group formation obtained with the plurality voting strategy. The item ratings involved in the first step of the algorithm are coloured.

	Item									
User	i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	10	4	3	6	10	9	6	8	10	8
u_2	1	9	8	9	7	9	6	9	3	8
u_3	10	5	2	7	9	8	5	6	7	6
group	**5**	**3**	**1**	**8**	**10**	**9**	**2**	**7**	**4**	**6**

Fig. 8. Group choice selection following the plurality voting strategy. The ranked list of items for the group could be (i_5, i_6, i_4, i_8, i_{10}, i_1, i_9, i_2, i_7, i_3), following the user selecting order u_1, u_2 and u_3, and setting *L*=3.

- **Approval voting strategy**. A threshold is considered for the item ratings: only those ratings greater or equal than the threshold value are taking into account for the profile combination. An item receives a vote for each user profile that has its rating surpassing the established threshold. The larger the number of votes the more influential the item is for the group (Figure 9). This strategy intends to promote the election of moderate alternatives: those that are not strongly disliked.

	Item									
User	i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	10	4	3	6	10	9	6	8	10	8
u_2	1	9	8	9	7	9	6	9	3	8
u_3	10	5	2	7	9	8	5	6	7	6

⇩ *threshold* = 5

	Item									
User	i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	1			1	1	1	1	1	1	1
u_2		1	1	1	1	1	1	1		1
u_3	1			1	1	1		1	1	1
group	**2**	**1**	**1**	**3**	**3**	**3**	**2**	**3**	**2**	**3**

Fig. 9. Group choice selection following the approval voting strategy. The ranked list of items for the group would be (i_4-i_5-i_6-i_8-i_{10}, i_1-i_7-i_9, i_2-i_3).

- **Borda count strategy** [6]. Scores are assigned to the items according to their ratings in a user profile: those with the lowest value get zero scores, the next one up one point, and so on. When an individual has multiple preferences with the same rating, the averaged sum of their hypothetical scores are equally distributed to the involved items. With the obtained scores, an additive strategy is followed, and the larger the sum the more influential the item is for the group. Figure 10 shows an example of the two steps followed by Borda count strategy. In the first step, ratings are normalised according to their relative relevance within the users' preferences. The items with the three lowest ratings for user u_1 are coloured in the tables. For the first one (in increasing rating value), d_3, a zero score is assigned. The second one, d_2, receives a score of value 1. The next score to be assigned would be 2. In this case, the next two items with lowest rating value, d_4 and d_7, have the same rating. In this case, two scores (2 and 3) are considered, and the average of them, i.e., (2+3)/2=2.5, is assigned to both items.

	Item									
User	i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	10	4	3	6	10	9	6	8	10	8
u_2	1	9	8	9	7	9	6	9	3	8
u_3	10	5	2	7	9	8	5	6	7	6

⇩

	Item									
User	i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	8	1	0	2.5	8	6	2.5	4.5	8	4.5
u_2	0	7.5	4.5	7.5	3	7.5	2	7.5	1	4.5
u_3	9	1.5	0	5.5	8	7	1.5	3.5	5.5	3.5
group	**17**	**10**	**4.5**	**15.5**	**19**	**20.5**	**6**	**15.5**	**14.5**	**12.5**

Fig. 10. Group choice selection following the Borda count strategy. The ranked list of items for the group would be (i_6, i_5, i_1, i_4-i_8, i_9, i_{10}, i_2, i_7, i_3).

- **Copeland rule strategy** [9]. Being a form of majority voting, this strategy sorts the items according to their *Copeland index*: the difference between the number of times an item beats (has higher ratings) the rest of the items and the number of times it loses. Figure 11 shows an example of Copeland rule strategy. In the bottom table, a +/– symbol in the *ij*-th cell (*i* for rows, and *j* for columns) means that item at *j*-th column was rated higher/lower than item at *i*-th row by the majority of the users. A zero value in a cell means that the corresponding items were rated with the same number of "beats" and "looses".

	Item									
User	i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	10	4	3	6	10	9	6	8	10	8
u_2	1	9	8	9	7	9	6	9	3	8
u_3	10	5	2	7	9	8	5	6	7	6

⇩

	Item									
User	i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
u_1	0	-	-	-	0	-	-	-	0	-
u_2	+	0	-	+	+	+	0	+	+	+
u_3	+	+	0	+	+	+	+	+	+	+
u_4	+	-	-	0	+	+	-	0	0	-
u_5	0	-	-	-	0	-	-	-	-	-
u_6	+	-	-	-	+	0	-	-	-	-
u_7	+	0	-	+	+	+	0	+	+	+
u_8	+	-	-	0	+	+	-	0	+	-
u_9	0	-	-	0	+	+	-	-	0	-
u_{10}	+	-	-	+	+	+	-	+	+	0
group	**+7**	**-6**	**-9**	**+1**	**+8**	**+5**	**-6**	**0**	**+3**	**-3**

Fig. 11. Group choice selection following the Copeland rule strategy. The ranked list of items for the group would be (i_5, i_1, i_6, i_9, i_4, i_8, i_{10}, i_2, i_7, i_3).

3 Group Recommender Systems

As stated by several authors [4, 7, 27], group recommender systems can be classified into two main categories: *aggregated models*, which aggregate individual user data into a group data, and generate predictions based on the group data; and *aggregated predictions*, which aggregate the predictions for individual users into group predictions. Other authors [10] have considered the way in which individual preferences are obtained (by content-based or collaborative filtering) as an additional dimension to be taken into account in such categorisation. In any of the above cases, the mechanisms in which user profile models or item predictions are aggregated are manifold, and can be based on any of the social choice strategies explained in Section 2.

In this section, we revise state of the art group recommendation approaches based on user model aggregation, and approaches based on prediction aggregation. We also briefly discuss approaches according to how groups are formed, and approaches that incorporate cooperative consensus to achieve a final recommendation policy agreed by the different members of a group.

3.1 Group Recommendation based on Model Aggregation

The group modelling problem has been addressed by **merging similar individual user profiles**. In this scenario, user profiles are usually represented as sets of weighted preferences or as sets of personal scores assigned by the users to the existing items.

INTRIGUE [2] is a tourist information server that presents information about the area around Torino, Italy. The system recommends sightseeing destinations and itineraries by taking into account the preferences of heterogeneous tourist groups, explains the recommendations by addressing the group members' requirements, and provides an interactive agenda for scheduling a tour. For each individual attraction, a record in a database stores characteristics and properties as a set of feature/value pairs, some of them related to geographical information and others used for matching preferences and interests of the users. Group recommendations are conducted in three steps. Firstly, the group is modelled as a set partitioned into a number of homogeneous subgroups, whose members have similar characteristics and preferences, and are assigned different degrees of influence on the estimation of the group preferences. Next, items are separately ranked by taking the preferences of each subgroup into account. Finally, subgroup-related rankings are merged to obtain the ranking suitable for the whole group.

In [17], Masthoff discusses several strategies based on social choice theory for merging individual user models to adapt to groups (see Section 2). Considering a list of TV programs, a group of viewers represent their interests with sets of personal 1-10 rating for the different TV programs. The author investigates how humans select a sequence of items for the group to watch, how satisfied people believe they would be with the sequence chosen by the different strategies, and how their satisfactions correspond with that predicted by a number of satisfaction functions. These evaluation functions are modified in [18], where satisfaction is modelled as a mood, and assimilation and decline of emotions with time is incorporated. Conducting a user study, she found that participants cared about fairness, and about preventing misery and starvation, as done in strategies like Average, Average without Misery, and Least Misery.

A more sophisticated strategy to merge various individual user profiles based on total distance minimisation is presented in [36]. The authors present a TV program recommender system for multiple viewers, in which the minimisation of the total distance between user profiles guarantees that the merged result could be close to most users' preferences. The shown experimental results prove that the resultant group profile actually reflects most members' preferences of the group.

An evaluation of profile aggregation strategies on a real large-scale dataset of TV viewings is presented in [27], showing that consensus-based strategies (especially the Utilitarian/Average strategy) provided the best recommendation results by comparing the built group profiles to a reference group profile obtained by directly analysing group consumptions.

In [7], we present an approach to automatically identify communities of interest from the tastes and preferences expressed by users in personal ontology-based profiles. The proposed strategy clusters those semantic profile components shared by the users, and according to the found clusters, several layers of interest networks are built. The social relations of these networks are finally used to provide group-oriented recommendations. In this context, we evaluate our approach by using different social choice strategies and, similarly to Senot and colleagues [27], found that consensus-based approaches outperformed borderline strategies, such as Least Misery, Most Pleasure and Plurality Voting strategies.

3.2 Group Recommendation based on Prediction Aggregation

In addition to group modelling, there exist several approaches that have been applied to the problem of making recommendations for groups of people under the framework of **aggregating lists of recommendations** for individual users belonging to the group. For them, we can distinguish two main strategies, namely *collaborative filtering* and *rank aggregation*.

In *collaborative filtering*, a user provides ratings to items, and these ratings are used to suggest her ranked lists that contain other items according to the overall preferences of people with similar rating patterns. Similarity rating patterns are calculated by using different metrics, such as Pearson and Spearman's correlations, and cosine-based distance.

In [14], a video recommender system is presented. Under a client/server architecture, the system receives and sends emails to obtain user ratings, and to provide video suggestions. The recommendations are shown to the users sorted by predicted ratings, and classified by video categories. The system also provides ranked lists from the most similar users, giving thus recommendations to a group of users (virtual community), instead of to individual users. The authors obtained open ended feedback from users indicating interest in establishing direct social contacts within their virtual community.

PolyLens [23] is a collaborative filtering system that suggests movies to groups of people with similar interests, which are expressed through personal five-start scale ratings from the well-known *MovieLens* recommender system [13]. In *PolyLens*, groups of people are explicitly created by users. For each member of a group, a ranked list of movies is obtained from a classic collaborative filtering mechanism. The individual ranked lists are merged according to the least misery principle, i.e., using a social value function where the group's happiness is the minimum of the individual members' happiness scores. Experimenting with *PolyLens*, the authors analysed primary design issues for group recommenders, such as the nature of the groups (in terms of persistency and privacy), the rights of group members, the social value functions for groups, and the interfaces for displaying group recommendations. They found that users not only valued group recommendations, but also were willing to yield some privacy to get the benefits of such recommendations, and extend the recommender system to enable them to invite non-members to participate, via email.

In *rank aggregation*, item recommendation lists are generated for each individual, and afterwards are merged into a single recommendation list for the group. Analogously to model aggregation approaches, different social choice strategies can be used to combine several rankings.

By exploring rank aggregation techniques on *MovieLens* dataset, Baltrunas and colleagues [4] showed that the effectiveness of group does not necessarily decrease when the group size groups, especially if the group have similar minded users. Moreover, they found that if individual recommendations are not correctly ranked (i.e. are not good enough), then recommending items ordered for a group can improve the effectiveness of individual recommendations.

This last result was also presented in [5]. The authors empirically evaluated a number of model aggregation and rank-based prediction aggregation techniques. By using a dataset of explicit ratings for recipes, provided by families of users in an e-health portal, they observed that (1) aggregating individual user models was superior to aggregating individual recommendations, and (2) role-based weighting outperformed uniform weighting.

3.3 Group Formation

Many studies have examined systems that support group formation. The groups can be built **intentionally** (by explicit definition from the users) or **non-intentionally** (by automatic identification from the system).

Kansas [28] is a virtual world in which *a user can explicitly join a group* by moving towards other users, who share a specific virtual spatial region to work collaboratively in a common task. Inside a group, the users can play different roles according to their current capabilities, which are defined by system treatments of user inputs and outputs. These capabilities can be manually acquired and dropped, or can be transferred by one user to another. The authors explain how direct manipulation and control, the “desktop metaphor”, might be an interesting approach for human computer interaction in cooperative environments.

MusicFX [21] enables *automatic group formation* by selecting music in a corporate gym according to the musical preferences of people working out at a given time. Thus, performing as a group preference arbitration system, *MusicFX* allows users to influence, but not directly control, the selection of music in the fitness centre. Specifically, each user specifies his preference for each musical genre. An individual preference rating for a genre is presented by a number ranging from –2 to +2. The group preference for that genre is then computed by the sum of the current users’ individual preferences. The system uses a weighted random selection policy for selecting one of the group top *N* music genres. One interesting anecdote the authors found with the system was the fact that people began modifying their workout times to arrive at the gym with other people, often strangers, who shared their music tastes.

3.4 Cooperative Consensus

In addition to applying an automatic group modelling algorithm, there exist approaches that make use of **consensus mechanisms** to achieve a final item recommendation policy agreed by the different members of a group. Recently, these approaches have also been called *role-based* [5] and *borderline* [27] strategies.

Travel Decision Forum [15] proposes a manual user interest aggregation method for group modelling by (1) allowing the current member optionally to view (and perhaps copy) the preferences already specified by other members, and (2) mediating user negotiations offering the users proposals and adaptations of their preferences. This method has several advantages, such as saving of effort, learning from other members, and encouraging assimilation to facilitate the reaching of agreement. In this system, neither user profile merging nor recommendation is used.

Collaborative Advisory Travel System, *CATS* [22], is a cooperative group travel recommender system which aims to help a group of users arrive at a consensus when they need to plan skiing holidays together; each having their own needs and preferences with respect to what constitutes as an ideal holiday for them. *CATS* system makes use of visual cues to create emphasis and help users locate relevant information, as well as enhance group awareness of each other’s preferences and motivational orientations. Individual user models are defined as set of critiques, i.e., restrictions on vacation features that should be satisfied. The system constructs a reliable group-preference model measuring the quality of each vacation package in terms of its compatibility with the restrictions declared by the members of the group.

4 Open Research Problems in Group Recommender Systems

Group recommender systems are still a novel research area. There are many challenging problems for further investigation. Masthoff has recently compiled some of such problems in [19]. Here, we summarise them and include others:

- *Dealing with uncertainty and scarcity in user profiles*. Issues like uncertain, non precise user preferences [10], and cold-start situations may also affect the accuracy of recommendations for certain members of a group.
- *Dealing with social dynamics in a group*. Members of a group may have complex social relationships (e.g., distinct roles, compromises, moods, ages) that affect the individuals' satisfaction for group recommendations. Multi-criteria and constrained recommenders may play a key role in such scenario.
- *Recommending item sequences to a group*. Already suggested items may influence the group members' satisfaction with subsequent recommendations.
- *Explaining recommendations to a group*. Showing how satisfied other members of the group are may improve the user' understanding of received recommendations, and may help to make her accepting suggestions of items she does not like. In this context, however, such transparency has to be balanced with aspects like privacy (e.g. to avoid the embarrassment effect) and scrutability. The reader is referenced to [34] for a detailed explanation of the roles of explanations in recommender systems.
- *Incorporating negotiation mechanisms*. Encouraging and supporting cooperation is a key aspect in many recommender systems. Facilitating a group of users to easily and friendly negotiate a final decision among a set of item recommendations may increase the individuals' satisfaction.
- *Designing user interfaces*. The user interface of a recommender may affect an individual's satisfaction with group recommendations. For example, in a TV show recommendation scenario, showing the current and the next items to be watched could increase the satisfaction of a user who does like the current suggested item, but is really keen on the subsequent one.
- *Evaluating group recommendations*. Better validation of satisfaction functions should be performed. Among other issues, large-scale evaluations [27], and studies on the affect of group size and composition (e.g., diversity of individuals' preferences within a group) have to be conducted [4, 5]. The reader is referenced to [26] for a detailed discussion of evaluation metrics and methodologies for recommender systems.

5 Group Recommender Systems for the Social Web

The Social Web is attracting millions of users, who are no longer mere consumers, but also producers of content. Social systems encourage interaction between users and both online content and other users, thus generating new sources of knowledge for recommender systems. The Social Web presents thus new challenges for recommender systems [12]. In the context of group recommendations, we can highlight the following research directions:

- *Developing new applications*. The huge amount and diversity of user generated content available in the Social Web allow investigating scenarios in which a group of individuals is recommended with "social objects" such as photos, music tracks and video clips stored in online multimedia sharing sites; stories, opinions and reviews published in blogs; and like-minded people registered in online social networks. In such applications, user generated content like ratings, tags, posts, personal bookmarks and social contacts could be exploited by novel group recommendation algorithms [12].
- *Dealing with dynamics and diversity of virtual communities*. In online social networks, people tend to reproduce or extend their relations in the real world to the virtual worlds conformed by the social networks. In [32], the authors show that relationship strength can be accurately inferred from models based on profile similarity and interaction activity on online social networks. Based on these findings, group recommender systems could incorporate content and social interests of group members to perform more accurate item suggestions. For such purpose, it would be necessary to investigate large group characteristics that impact individual decisions, and explore new satisfaction and consensus functions that capture social, interest, and expertise (dis)similarity among the members of a community [11]. With this respect, because of the evolving composition of online communities, analysing and exploiting the time dimension in the above characteristics may play a key role to obtain more accurate recommendations for community members.
- *Incorporating contextual information*. The anytime-anywhere phenomenon is present in any social system and thus, group recommenders for the Social Web should incorporate contextual information [1]. They would have to automatically detect user presence from inputs provided by mobile, sensor and social data sources [32], and adaptively infer the strength of the social connections within the group, in order to provide accurate recommendations.
- *Finding communities of interest*. In the Social Web, it is very often the case that the membership to a community is unknown or unconscious. In many social applications, a person describes her interests and knowledge in a personal profile to find people with similar ones, but she is not aware of the existence of other (directly or indirectly) related interests and knowledge that may be useful to find those people. Furthermore, depending on the context of application, a user can be interested in different topics and groups of people. In both cases, for individual and group recommender systems, a strategy to automatically identify communities of interest could be very beneficial [7].
- *Integrating user profiles from multiple social systems*. Increasingly, users maintain personal profiles in more and more Web 2.0 systems, such as social networking, personal bookmarking, collaborative tagging, and multimedia

sharing sites. Recent studies have shown that inter-linked distributed user preferences expressed in several systems not only tend to overlap, but also enrich individual profiles [31, 35]. A challenging problem in the recommender system field is the issue of integrating such sources of user preference information in order to provide the so called cross-domain recommendations [35]. This clearly opens new research opportunities for group recommenders, which e.g. could suggest to a virtual community sharing interests in a particular domain with items belonging to other domain but liked by some of its members, e.g. recommending specific pieces of classical music to a group with interests in 18th century art.

The authors of this chapter have explored some of the above research paths. We have investigated the use of explicit semantic information as an enhanced modelling ground to combine individual user preferences [30]. We have also researched methods to find implicit communities of interest as a form of latent groups, by mining the kind of user input that is commonly available to a recommender system, along with additional semantic data [7]. We have found an inverse role to the usual one for user communities and groups: besides their natural purpose as user aggregation units, groups provide a basis for user model decomposition. We investigated the use of group models as projecting spaces, to produce sections of user interests –subprofiles– by a projection of complete profiles into the subspace induced by the group model. We found that subprofiles enable more focused and in some situations more precise recommendations than user profiles treated as indivisible units.

6 Conclusions

With the advent of the Social Web, people more and more often join virtual communities and social networks, and participate in many different types of collaborative systems, such as wiki-style, product reviewing, and multimedia sharing sites, among others. This together with the progressive spreading of ambient intelligence technologies (e.g., location and mobile-based sensors) in open environments bring in new appealing possibilities and problems for the recommender systems research agenda, which are related to suggesting interesting "social objects" (multimedia items, people, events, plans, etc.) to groups of people having explicit or implicitly bonds among them.

In this chapter, we have revisited existing approaches to group recommendations, and have discussed open research problems in the area, extending such discussion towards a number of potential new research directions related to the context of the Social Web. New complexities and compelling perspectives emerge for recommender systems oriented to groups of users of very different nature and size, as the ones currently growing in the Social Web.

Acknowledgements

This work was supported by the Spanish Ministry of Science and Innovation (TIN2008-06566-C04-02), and the Community of Madrid (S2009TIC-1542).

References

1. G. Adomavicius, A. Tuzhilin. Context-Aware Recommender Systems. In *Recommender Systems Handbook*, F. Ricci, L. Rokach, B. Shapira, P. B. Kantor (Eds.), pp. 217--253 (2011)
2. L. Ardissono, A. Goy, G. Petrone, M. Segnan, P. Torasso. INTRIGUE: Personalized Recommendation of Tourism Attractions for Desktop and Handset Devices. *Applied Artificial Intelligence* **17**(8-9), pp. 687--714 (2003)
3. R. Baeza-Yates, B. Ribeiro Neto. Modern Information Retrieval. Addison-Wesley (1999)
4. L. Baltrunas, T. Makcinskas, F. Ricci. Group Recommendations with Rank Aggregation and Collaborative Filtering. In *Proceedings of the 4th ACM Conference on Recommender Systems (RecSys 2010)*, pp. 119--126 (2010)
5. S. Berkovsky, J. Freyne. Group-based Recipe Recommendations: Analysis of Data Aggregation Strategies. In *Proceedings of the 4th ACM Conference on Recommender Systems (RecSys 2010)*, pp. 111--118 (2010)
6. J. C. Borda. Mémoire sur les Élections au Scrutin. Histoire de l'Académie Royale des Sciences (1781)
7. I. Cantador, P. Castells. Extracting Multilayered Communities of Interest from Semantic User Profiles: Application to Group Modeling and Hybrid Recommendations. *Computers in Human Behavior*. Elsevier. In press (2011)
8. D. L. Chao, J. Balthrop, S. Forrest. Adaptive Radio: Achieving Consensus using Negative Preferences. In *Proceedings of the 2005 International ACM Conference on Supporting Group Work* (GROUP 2005), pp. 120--123 (2005)
9. A. H. Copeland. A Reasonable Social Welfare Function. Seminar on Applications of Mathematics to the Social Sciences, University of Michigan (1951)
10. L. M. De Campos, J. M. Fernández-Luna, J. F. Huete, M. A. Rueda-Morales. Managing Uncertainty in Group Recommending Processes. *User Modeling and User-Adapted Interaction* **19**(3), pp. 207--242 (2009)
11. M. Gartrell, X. Xing, Q. Lv, A. Beach, R. Han, S. Mishra, K. Seada. Enhancing Group Recommendation by Incorporating Social Relationship Interactions. In *Proceedings of the 16th ACM International Conference on Supporting Group Work (GROUP 2010)*, pp. 97--106 (2010)
12. W. Geyer, J. Freyne, B. Mobasher, S. S. Anand, C. Dugan. 2nd Workshop on Recommender Systems and the Social Web. In *Proceedings of the 4th ACM Conference on Recommender Systems (RecSys 2010)*, pp. 379--380 (2010)
13. J. Herlocker, J. A. Konstan, A. Borchers, J. Riedl. An Algorithmic Framework for Performing Collaborative Filtering. In *Proceedings of the 22nd ACM Conference on Research and Development in Information Retrieval (SIGIR 1999)*, pp. 230--237 (1999)

14. W. Hill, L. Stead, M. Rosenstein, G. Furnas. Recommending and Evaluating Choices in a Virtual Community of Use. In *Proceedings of the 13th International Conference on Human Factors in Computing Systems (CHI 1995)*, pp. 194--201 (1995)
15. A. Jameson. More than the Sum of its Members: Challenges for Group Recommender Systems. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI 2004)*, pp. 48--54 (2004)
16. A. Jameson, B. Smyth. Recommendation to Groups. In *The Adaptive Web*, P. Brusilovsky, A. Kobsa, W. Nejdl (Eds.), pp. 596--627 (2007)
17. J. Masthoff. Group Modeling: Selecting a Sequence of Television Items to Suit a Group of Viewers. *User Modeling and User-Adapted Interaction* **14**(1), pp. 37--85 (2004)
18. J. Masthoff. The Pursuit of Satisfaction: Affective State in Group Recommender Systems. In *Proceedings of the 10th International Conference on User Modeling (UM 2005)*, pp. 297--306 (2005)
19. J. Masthoff. Group Recommender Systems: Combining Individual Models. In *Recommender Systems Handbook*, F. Ricci, L. Rokach, B. Shapira, P. B. Kantor (Eds.), pp. 677--702 (2011)
20. J. F. McCarthy. Pocket RestaurantFinder: A Situated Recommender System for Groups. In *Proceedings of the ACM CHI 2002 International Workshop on Mobile Ad-Hoc Communication* (2002)
21. J. F. McCarthy, T. D. Anagnost. MusicFX: An Arbiter of Group Preferences for Computer Supported Collaborative Workouts. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work (CSCW 1998)*, pp. 363--372 (1998)
22. K. McCarthy, M. Salamo, L. McGinty, B. Smyth. CATS: A Synchronous Approach to Collaborative Group Recommendation. In *Proceedings of the 19th International Florida Artificial Intelligence Research Society Conference (FLAIRS 2006)*, pp. 1--16 (2006)
23. M. O'Connor, D. Cosley, J. A. Konstan, J. Riedl. PolyLens: A Recommender System for Groups of Users. In *Proceedings of the 7th European Conference on Computer Supported Cooperative Work (ECSCW 2001)*, pp. 199--218 (2001)
24. P. K. Pattanaik. Voting and Collective Choice. Cambridge University Press (2001)
25. S. Pizzutilo, B. De Carolis, G. Cozzolongo, F. Ambruoso. Group Modeling in a Public Space: Methods, Techniques, Experiences. In *Proceedings of the 5th WSEAS International Conference on Applied Informatics and Communications (AIC 2005)*, pp. 175--180 (2005)
26. G. Shani, A. Gunawardana. Evaluating Recommendation Systems. In *Recommender Systems Handbook*, F. Ricci, L. Rokach, B. Shapira, P. B. Kantor (Eds.), pp. 257--297 (2011)
27. C. Senot, D. Kostadinov, M. Bouzid, J. Picault, A. Aghasaryan, C. Bernier. Analysis of Strategies for Building Group Profiles. In *Proceedings of the 18th International Conference on User Modeling, Adaptation, and Personalization (UMAP 2010)*, pp. 40--51 (2010)
28. R. B. Smith, R. Hixon, B. Horan. Supporting Flexible Roles in a Shared Space, In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work (CSCW 1998)*, pp. 197--206 (1998)
29. B. Smyth, E. Balfe, J. Freyne, P. Briggs, M. Coyle, O. Boydell. Exploiting Query Repetition and Regularity in an Adaptive Community-Based Web Search Engine. *User Modeling and User-Adapted Interaction* **14**(5), pp. 383--423 (2005)

30. M. Szomszor, H. Alani, I. Cantador, K. O'Hara, N. R. Shadbolt. Semantic Modelling of User Interests based on Cross-Folksonomy Analysis. In *Proceedings of the 7th International Semantic Web Conference (ISWC 2008)*, pp. 632--648 (2008)
31. M. Szomszor, I.Cantador, H. Alani. *Correlating User Profiles from Multiple Folksonomies*. In *Proceedings of the 19th ACM Conference on Hypertext and Hypermedia (Hypertext 2008)*, pp. 33--42 (2008)
32. M. Szomszor, C. Cattuto, W. Van den Broeck, A. Barrat, H. Alani. Semantics, Sensors, and the Social Web: The Live Social Semantics Experiments. In *Proceedings of the 7th Extended Semantic Web Conference (ESWC 2010)*, vol. 2, pp. 196--210 (2010)
33. A. D. Taylor. Mathematics and Politics: Strategy, Voting, Power and Proof (1995)
34. N. Tintarev, J. Masthoff. Designing and Evaluating Explanations for Recommender Systems. In *Recommender Systems Handbook*, F. Ricci, L. Rokach, B. Shapira, P. B. Kantor (Eds.), pp. 479--510 (2011)
35. P. Winoto, T. Ya Tang. *If You Like the Devil Wears Prada the Book, Will You also Enjoy the Devil Wears Prada the Movie? A Study of Cross-Domain Recommendations*. New Generation Computing **26**(3): 209--225 (2008)
36. Z. Yu, X. Zhou, Y. Hao, J. Gu. TV Program Recommendation for Multiple Viewers Based on user Profile Merging. *User Modeling and User-Adapted Interaction* **16**(1), 63--82 (2006)