Group Recommendation for Smart Applications: a Multi-agent View of the Problem

Francesco Barile
Dipartimento di Fisica,
Università degli Studi di Napoli
“Federico II”, via Cintia MSA,
80126 Napoli, Italy
francesco.barile@unina.it

Antonio Caso
Dipartimento di Fisica,
Università degli Studi di Napoli
“Federico II”, via Cintia MSA,
80126 Napoli, Italy
antonio.caso@unina.it

Silvia Rossi
Dipartimento di Ingegneria Elettrica
e Tecnologie dell'Informazione,
Università degli Studi di Napoli
“Federico II”, via Claudio 21,
80125 Napoli, Italy
silvia.rossi@unina.it

***Abstract*—The widespread use of social networks is modifying our way to live our cities and to plan/decide our activities. The long-term goal of our research is to provide users with group recommendation and decision support systems for smart city applications that rely on the analysis of the users' behaviors on social networks/media. In this paper, we provide an overview of the group recommendation literature from a game-theoretical/multi-agent system classical point of view, and we present a practical example where a social choice mechanism is extended with social information extracted from the analysis of the interactions, within small groups of users, on a social network.**

I. Introduction

The widespread use of social networks is modifying our way to live our cities and to plan/decide our activities. In particular, the way we plan a travel, select places to visit or choose a restaurant, is influenced by the information we can access on social media. In this direction, the long-term goal of our research is to provide users with recommendation and decision support systems for smart city applications that rely on the analysis of the users' behaviors on social networks/media. Examples of such applications are city tour planners or activities recommendation systems. Moreover, an automatic outdoor planning system of a city tour has to take into account that, potentially, groups of users (and not a single user) jointly select the activities to perform. Hence, it is necessary to choose activities, or, more generally speaking, certain Points of Interest (POI) that maximize the group satisfaction, considering that the members' preferences can be different.

The problem of group recommendation has been widely studied in the fields of information retrieval, mathematics, economics and multi-agent systems. Group recommendation approaches rely either on building a single group profile, resulting from the combination of the profiles of all the users, or on merging the recommendation lists generated for each individual users, at runtime, using different group decision strategies. Nevertheless, many of these techniques do not consider social relationships among the group members [1], while the design and implementation of group recommendation systems, and, more generally, of decision support systems, should take into account the type of control in the group decision-making process [2]. PolyLens [3] has been one of the first approaches to include social characteristics within a group recommendation system. A more recent example is represented by the work of [1], where the Authors started to evaluate the group members' weights, in terms of individual members' importance or influence in a group, for movie recommendations.

In this work, we provide a general overview on the literature on group recommendations from a multi-agent/game theoretical point of view. Finally, we present a practical example where a classical social choice mechanism is extended with social information extracted from the analysis of the interactions, among members of small groups, in a social network. Differences between the classical implementation of the function and its extended version are evaluated with a user study.

II. Problem Definition

In our work, we are interested in providing recommendations to groups of friends, intended as sets of people that meet, interact, or have some actual common bond in the physical world, since they are planning real activities to perform together. Moreover, the identity of the group members has to be dynamically determined, since the actual members of a group can be established only according to the activity to perform (i.e., a travel, a movie, and so on).

Following these requirements, we decided to focus on group recommendation systems that aggregate the recommendations considered for the single users. This approach will provide us the flexibility required in the group formation process (single user's profile and recommendations are build independently from the group's members) and to dynamically account group relationships, at the time of providing the group recommendations (the users' recommendations are merged only once the group is formed). Finally, our system will not have to provide a single item recommendation, but sequences or sets, as required by the touristic application domain.

More formally, given a set of n friends ($F = \{1, \dots, n\}$) and a set of m POI ($P = \{1, \dots, m\}$), each user $i \in F$ has a preference profile $\succ_i$ over P ($\succ_i = \{r_i(1), \dots, r_i(m)\}$) with $r_i(x) \in \mathcal{R}$, which represents the user i rank for the x POI, and is is evaluated by a single user recommendation agent. Our goal is to design a mechanism to obtain $\succ_F = \{r_F(1), \dots, r_F(m)\}$, where $r_F(x)$ is the correspondent ranking for the x POI, as evaluated for the group, and that will help the group decision by simplifying the consensus making process.

III. A Multi-agent View

In this section, we provide an overview on the multi-agent system (MAS) literature dealing with the problem of group recommendation. In this context, we can model the members of the group as agents, and they interact together to find the best compromise for the whole group. There are several approaches proposed in literature, which include classical game theory, voting, coalition making, social choice analysis and negotiation.

A. Social Choice

A typical approach for merging user rankings is the definition of Social Choice functions. Social Choice strategies, according to [4], can be classified as majority-based strategies (mainly implemented as voting mechanisms to determine the most popular choices among alternatives), consensus-based strategies (that try to average among all the possible choices and preferences), and role-based strategies (that explicitly takes into account possible roles and hierarchical relationships among members). In the rest of this section, we briefly summarize the main strategies for each category, providing references for further details.

a) Average Strategy: as described in [5], [1], this approach consists in computing the group rating for an activity as the average of all group members' ratings. It is commonly used as a benchmark for comparison or as base to define more complex approaches.

b) Fairness Strategy: this strategy, described in [6], [5], needs of an ordering among the users of the group. In the simplest case this ordering can be random. The first user is selected and his k best rated activities are taken into account. From them, the activity that guarantees the less misery to the other group members is chosen. The process is iterated on the other users until k activities are selected.

c) Borda Count Strategy: this strategy, as explained in [6], [7], [8], consists into two phases. Initially, users' ratings are replaced with scores, assigned in this way: the activity with the lowest rate for a user gets zero score, the next one gets one, and so on. If two or more activities have the same rating they are assigned with the average of the scores that should have. After that, an additive or an average strategy can be used on those scores to obtain the group recommendation.

d) Plurarity Voting Strategy: as for the fairness strategy, also in this case a sorting among users is necessary. Starting from the first user, his/her k top rated activities are considered. In this case, however, among these the activity more voted by other users is selected. The strategy is explained in [6], [8].

e) Least Misery Strategy: this strategy, used in [3] for a group recommendation system for movie, *PolyLens*, can be used when one or more users give a rating particularly low for some activities. In case of small groups, it is reasonable to assume that the satisfaction of the group that performs an activity could decrease if one or more components really dislike the activity. Least Misery Strategy consists in assigning to each activity the minimum of its ratings.

f) Most Pleasure Strategy: if some user really likes one activity that is acceptable for other group members it should be rational to use, as group rating, the greatest given rating for the activity.

B. Negotiation

Another way, proposed in literature, to address the group recommendation problem, is the use of Negotiation methodology. In general, a set of agents act on behalf of human group members, participating in a cooperative negotiation for generating the group recommendations.

In [9] there is a negotiation agent for each group member. An individual recommendation system gives recommendation for a set of items, and, in addition to this, an individual utility of each product for each user is evaluated, introducing a user preference model. The Negotiation protocol is different according to the cardinality of the group. For groups of two people, the used protocol is of *alternating offers*, while for groups of more people, it uses a *merging ranks* protocol, with a mediator agent that has strategies to help in choosing among proposals and by offering an agreement to the group (i.e, by maximizing the average utilities of group members or maximizing the utility of the least happy member). The framework is tested by simulating the negotiation protocols.

Another approach is proposed in [10], where an *alternating offers* protocol is used. In this approach there is not a mediation, but groups size is restricted at two users. There is an agent for each user, and there is a two-level user profiling, which includes a *recommendation profile*, containing personal information and preferences, and a *negotiation profile*, used to distinguish agent behaviors in the negotiation among three degree of collaboration (self-interested, collaborative and highly collaborative). If the negotiation process finishes with an agreement among all the agents, the result is a list of constraints that match the preferences of the group members.

The original idea was then developed in a subsequent work [11], where user agents are configurable in order to exhibit the desired behavior of the corresponding user. The negotiation model is a multi-party negotiation that centralize the communications through a negotiator agent, acting as mediator. It receives the proposals of the user agents, combines them into a single proposal, which is later broadcasted by the negotiator agent and analyzed by the user agents. The system uses a domain ontology to describe the user's likes and the items to recommend. The user agent is responsible of building and updating a user profile, of obtaining the individual preference model, of participating in the negotiation process and of informing the user about the result of the negotiation. Besides, there are two support agents that help in computing the individual preference model (preferences agent) and in selecting the list of items that satisfy the group preferences, given the group preference model (items selector agent). The protocol used in the negotiation is a generalization of the bilateral alternating offers protocol [12] for the multi-party negotiation.

C. Coalitions

Differently from the previous cases, the use of coalitions in the field of group recommendation is not straightforward, since coalitions require the groups formation at run-time. The idea is to organize group members into smaller and cohesive groups, so it is possible to provide more effective recommendations to each of them. In [13] the problem is modeled as a coalitional game, where people are grouped into disjoint coalitions to maximize the social welfare function of the group. The payoff function considers the similarity between coalition members's ratings, and a weighting factor for the coalition size. The approach is compared with a classic K-Means clustering, on randomly formed groups, and the results shows better performances in the formation of larger coalitions. In some cases, however, this approach is not applicable, because it is not possible to reorganize the group into more cohesive sub-groups, but it is necessary to provide a recommendation for the whole group of users.

D. Normal Games

In some cases, group members could have more different interests conflicting with each other. In case of great heterogeneity, the attempt to resolve the conflict by applying a cooperative approach can lead to a failure in the negotiation [14]. In this scenario, it can be reasonable to apply non-cooperative approaches. The idea is that users can be viewed as self-interested agents and the recommendation system can be modeled as a classical non-cooperative game in normal form.

In [14] an alternative approach based on non-cooperative games is proposed. In this case, group members are viewed as the players of the game, the items to recommend are viewed as game actions, and the recommendation problem is modeled as a problem of finding the Nash Equilibrium for the game. This approach is compared with other state-of-art aggregation strategies, Average, Least misery (LM) and Plurality Voting (PV). Average strategy shows the best performance, but the proposed strategy performs better with respect to LM and PV when the groups become more heterogeneous and wider.

E. Weighted Utilities

In the previous sections, we presented the main MAS techniques used to address the group recommendation problem. The results presented in the literature showed that there is no strategy can be defined as the "best", but different approaches are better suited in different scenarios, depending from the characteristics of the specific group. Besides, traditionally MAS techniques do not seem to capture all the features of real-world scenarios. For example, automatic voting/ranking mechanisms often require that all the agents involved have the same influence on the decision procedure, while real group interactions take into account intra-group roles and mutual influences. Again, some members of the group could have a particular influence on the others, based on their personal experiences or on the strength of their mutual relationship. Furthermore, there may be situations where the participants follow a democratic process in order to find a possible solution, and cases where the group is supported by a human leader. Usually the decision of a group member whether or not to accept a given recommendation may depend not only on his/her own evaluation of the content of the recommendation, but also on his/her beliefs about the evaluations of the other group members [6]. Recommendation systems for groups need to capture both the preferences of the group members but also these key factors in the group decision process [1].

On the basis of these considerations it appears necessary to integrate information from the social relationships among group members with the classical MAS techniques and so to derive new strategies more applicable to the considered settings. The most common approach is to remodel the utility of the agents by applying weights derived from social interactions between the members of the group.

For example, the work of [1] starts to evaluate the group members weights, in terms of individual group members importance or influence for movie recommendations. The defined group consensus function relies on the concept of "expertise" and "group dissimilarity". Also, it introduces the idea of diversifying the social choice strategy to use on the basis of the characteristics of the group. The Authors calculate a "social value" on the basis of social interactions between group members, and then use this value to discriminate the strategy to apply for the group. Both expertise and social interaction values are derived from questionnaires.

Another way to use interactions between group members is presented in [15]. Here, the authors introduce the concept of *empathetic utility* on social networks: the satisfaction of an individual depends from both his intrinsic utility and his *empathetic utility* deriving from the happiness of his neighbors in the social network [15]. Based on this idea, individual preferences are aggregated in a weighted social choice function that takes into account local relationships with neighborhoods in the network. However, in [15] the Authors do not specify how to evaluate such numerical relationships, while they focus on computational aspects of scaling up with large networks of friends.

Both these approaches have the same gap, because they do not provide a way to automatically retrieve social information. An idea to address this problem is showed in the next section, where we illustrate a way to retrieve social information from Online Social Networks (OSNs).

IV. A LEADERSHIP WEIGHTED MODEL

In [16] we presented a simple "non semantic" approach to obtain information about users' leadership values in decision making, by analyzing the popularity of each user within the group. Such popularity values are obtained implementing an extension of the well-known PageRank algorithm [17] starting from the users' mutual interactions on the social network *facebook.com*. In fact, Online Social Networks analysis can provide a viable way to obtain, without intruding the users with questionnaires, information about the users and their social relationships within communities and groups. More in detail, this approach defines a centrality measure that takes into account the degree of activity of a person and the directionality of specific communication activities between pairs of users, using a combination of data collected from the OSN. Furthermore, as in the classic PageRank, each user inherits a portion of popularity from other users.

In this paper, we use such evaluation to weight an Average Satisfaction strategy. According to [18], users involved in real interaction seem to care about fairness and to avoid misery,

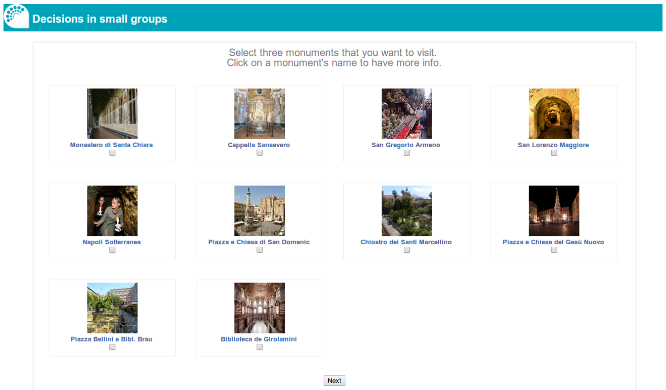

Fig. 1. A screen-shot of the web page used to select the activities to perform.

Fig. 2. A four people group taking the final decision.

while averaging among choices get good results. Inspired by the works of [1], [19], we defined a strategy that takes into account the leadership values as weights for the POI rankings, provided by the users (note that the sum of the leadership values among a group is equal to one). The proposed strategy to evaluate the group $r_F(x)$ rating for the POI x is the following:

$$r_{avg}(x) = \frac{1}{n}\sum_{i=1}^{n}(R(i) \cdot r_i(x)) \tag{1}$$

where n is the number of users in the group F, $R(i)$ is the leadership value of user i, calculated as in [16]. Hence, Equation 1 is a function that evaluates the average of all the i users rankings $r_i(x)$ weighted by the i-th leadership value $R(i)$.

The set $\succ_{avg} = \{r_{avg}(1), \ldots, r_{avg}(m)\}$, which is the set of group's rankings computed for each item, is then used to get the final decision: the first k activities x (with k equals to the number of activities to propose) with the higher $r_{avg}(x)$ values are selected for the recommendation. Moreover, in order to evaluate our function, we also implemented the standard version of a simple averaging function ($r_{st.avg}(x)$) on the same data:

$$r_{st.avg}(x) = \frac{1}{n}\sum_{i=1}^{n} r_i(x) \tag{2}$$

V. A Practical Example

We conducted a pilot study with real users involved in the task of planning a trip in a city. We evaluated only the users' behaviors in the decision making process and the impact of our consensus function on a simple case of binary selection of POI (i.e., yes or no decisions on a POI, without expressing an explicit ranking among activities). The behavior of 14 groups composed, in the average, of 3.36 close friends is evaluated. 46 users took part in the experimentation (26 men and 20 women). The average age was 27.3 with a graduate education.

We asked each user to register on a specific web site using the credentials of facebook.com. Once registered, he/she was asked to imagine to plan a one-day visit in a specific city and to select, from a checklist of ten items, only three activities (i.e., places to visit) for the day. After that, it was asked to select two restaurants (from a check list of eight). Since we do not want the users to be involved in strategic reasoning, we did not ask the users to express ratings and preferences among the selected choices. A screen-shot of the interface used to select the activities to perform is shown in Figure 1. For each $i \in F$ a vector $\succ_i = \{r_i(1), \ldots, r_i(m)\}$, with $r_i(x) \in \{0, 1\}$, and $m = 18$, representing the user choices, is stored, such that $\sum_{x \in P} r_i(x) = 5$.

The group was, then, asked to discuss, face-to-face, in order to obtain a shared and unique decision for the group. This final decision corresponds to the Ground Truth vector $\succ_{GT}$ used to evaluate our functions. Figure 2 shows a group while discussing the final choices with the support of a personal computer.

A. Results

The average number of analyzed users' interactions, within the entire group, is 1079 with a standard deviation of 1254. A very high standard deviation means that, for the specific class of users involved in the study (i.e., groups of close friends), on the considered OSN, the groups' (and members') behaviors were very different, and hence, they cover a wide range of possible users.

Firstly, the similarity of the proposed weighted version of

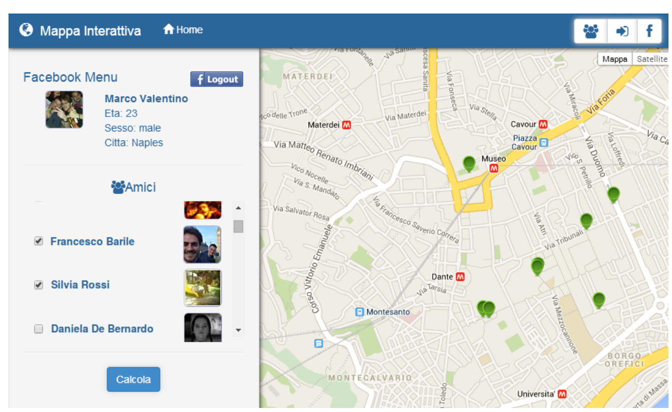

Fig. 3. An example of web application that implements the proposed approach.

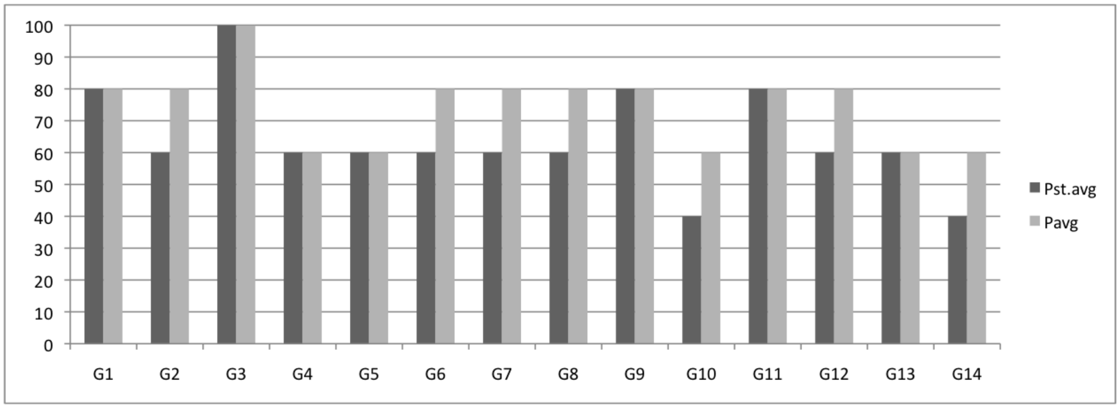

Fig. 4. $r_{st.avg}$ and r_{avg} results for each group.

% Similarity	$r_{st.avg}$	r_{avg}	r_{Leader}	r_i
mean	64 ± 16	74 ± 12	61 ± 17	59 ± 11

TABLE I. PERCENTAGES OF SIMILAIRTY RESULTS IN THE USER STUDY.

the average satisfaction function (r_{avg}) with respect to the groups' ground truth (r_{GT}) was evaluated. Such similarity is computed as a percentage of the r_{avg} choices that were already selected in the group final choices r_{GT}. We also evaluated the similarity of the groups' ground truth with respect to the standard implementation of such function (i.e., $r_{st.avg}$ as a typical averaging function on users' choices). Aggregated results are reported in Table I and, for each of the 14 groups, r_{avg} and $r_{st.avg}$ values are reported in Figure 4. With respect to their standard implementation, the function that takes into account social relationships perform slightly better (74% w.r.t. 64%). The r_{avg} consensus function often guesses 4 on 5 activities. In detail, as shown in Figure 4, taking into account the similarity of values for each group, the r_{avg} results are always better or equal with respect to $r_{st.avg}$ results.

Moreover, we evaluated the similarity of r_{GT} with respect to the leader's choices (r_{Leader}) of each group, and the similarity of r_{GT} with respect to the remaining users' choices in the group (r_i). The leader of each group was evaluated as the member with the highest $R(i)$ score according to leadership value calculated as indicated in [16]. Table I summarizes the cumulative data (mean values and standard deviations) of all groups involved in the experiment for such similarities. Considering the average similarity of each r_i with respect to the r_{GT}, the groups show a good value of cohesion (59%). Moreover, considering the aggregated data, the average similarity value of the leaders' choices is 61%, which is comparable with the r_i similarity value. This behavior is in accordance with our requirement to analyze close groups without any dictatorial user, or strong hierarchical relationships.

At last, we implemented the proposed approach to calculate POI recommendations for small groups of friends in a touristic web application. Figure 3 shows a user logged-in with facebook.com credentials and that selected two friends to create a group; the application, then, elaborates the global ranks for the group and filters the POI to show.

VI. Conclusion

The smarter cities paradigm relies on a massive use of technologies and big data analysis to enhance and facilitate human-environment interactions in a pervasive way. In order to be easily accessed and retrieved all these data has to be filtered and selected according to user behaviors and profiles. In particular, in this paper, we introduced the problem of group recommendation, as required in the developments of automatic and intelligent tools for activities scheduling.

The problem of group recommendation has been widely studied in different fields; here, we presented this issue as approached from MAS literature of view. Nevertheless, only few of the presented approaches started to consider social relationships among group members in the design of group recommendation, while almost all of these do not provide a mechanism to automatically retrieve this information.

Finally, in this paper, we presented a simple social choice function that uses "non-semantic" information extracted from real interactions on a social network to weight users' ratings/choices. We provided an evaluation of the strategy with respect to its standard implementation. Results showed that, even for very simple aggregation functions, the introduction of social relationships data might provide improvements in the recommendation process.

Acknowledgment

The research leading to these results has received funding from the Italian Ministry of University and Research and EU under the PON OR.C.HE.S.T.R.A. project (ORganization of Cultural HEritage for Smart Tourism and Real-time Accessibility).

References

[1] M. Gartrell, X. Xing, Q. Lv, A. Beach, R. Han, S. Mishra, and K. Seada, "Enhancing group recommendation by incorporating social relationship interactions," in *Proceedings of the 16th ACM International Conference on Supporting Group Work*, ser. GROUP '10. ACM, 2010, pp. 97–106.

[2] M. T. Jelassi and A. Foroughi, "Negotiation support systems: an overview of design issues and existing software," *Decision Support Systems*, vol. 5, no. 2, pp. 167 – 181, 1989.

[3] M. O'Connor, D. Cosley, J. A. Konstan, and J. Riedl, "Polylens: A recommender system for groups of users," in *Proceedings of the Seventh Conference on European Conference on Computer Supported Cooperative Work*, ser. ECSCW'01. Kluwer Academic Publishers, 2001, pp. 199–218.

[4] C. Senot, D. Kostadinov, M. Bouzid, J. Picault, A. Aghasaryan, and C. Bernier, "Analysis of strategies for building group profiles," in *User Modeling, Adaptation, and Personalization*, ser. Lecture Notes in Computer Science, P. Bra, A. Kobsa, and D. Chin, Eds. Springer Berlin Heidelberg, 2010, vol. 6075, pp. 40–51.

[5] J. Masthoff, "Modeling a group of television viewers," in *Proceedings of the Workshop Future tv, in Intelligent Tutoring Systems Conference*, 2002, pp. 34–42.

[6] I. Cantador and P. Castells, "Group recommender systems: New perspectives in the social web." in *Recommender Systems for the Social Web*, ser. Intelligent Systems Reference Library, J. J. P. Arias, A. F. Vilas, and R. P. D. Redondo, Eds. Springer, 2012, vol. 32, pp. 139–157.

[7] S. Ghosh, M. Mundhe, K. Hernandez, and S. Sen, "Voting for movies: The anatomy of a recommender system," in *Proceedings of the Third Annual Conference on Autonomous Agents*, ser. AGENTS '99. New York, NY, USA: ACM, 1999, pp. 434–435.

[8] J. Masthoff, "Group modeling: Selecting a sequence of television items to suit a group of viewers," in *Personalized Digital Television*. Springer, 2004, pp. 93–141.

[9] P. Bekkerman, S. Kraus, and F. Ricci, "Applying cooperative negotiation methodology to group recommendation problem," in *Proceedings of Workshop on Recommender Systems in 17th European Conference on Artificial Intelligence (ECAI 2006)*. Citeseer, 2006, pp. 72–75.

[10] I. Garcia, L. Sebastia, and E. Onaindia, "A negotiation approach for group recommendation." in *Proceedings of the 2009 International Conference on Artificial Intelligence, Las Vegas Nevada, USA*. CSREA Press, 2009, pp. 919–925.

[11] I. Garcia and L. Sebastia, "A negotiation framework for heterogeneous group recommendation," *Expert Systems with Applications*, vol. 41, no. 4, pp. 1245–1261, 2014.

[12] A. Rubinstein, "Perfect equilibrium in a bargaining model," *Econometrica: Journal of the Econometric Society*, pp. 97–109, 1982.

[13] L. A. M. Carvalho and H. T. Macedo, "Generation of coalition structures to provide proper groups' formation in group recommender systems," in *Proceedings of the 22Nd International Conference on World Wide Web Companion*, ser. WWW '13 Companion. International World Wide Web Conferences Steering Committee, 2013, pp. 945–950.

[14] L. A. M. C. Carvalho and H. T. Macedo, "Users' satisfaction in recommendation systems for groups: An approach based on noncooperative games," in *Proceedings of the 22Nd International Conference on World Wide Web Companion*, ser. WWW '13 Companion. International World Wide Web Conferences Steering Committee, 2013, pp. 951–958.

[15] A. Salehi-Abari and C. Boutilier, "Empathetic social choice on social networks," in *13th International Conference on Autonomous Agents and Multiagent Systems*, 2014, pp. 693–700.

[16] A. Caso and S. Rossi, "Users ranking in online social networks to support pois selection in small groups," in *Extended Proceedings of the 22nd Conference on User Modelling, Adaptation and Peronalization - UMAP 2014*, 2014, pp. 5–8.

[17] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Computer Networks and {ISDN} Systems*, vol. 30, no. 17, pp. 107 – 117, 1998, proceedings of the Seventh International World Wide Web Conference.

[18] J. Masthoff, "Group recommender systems: Combining individual models," in *Recommender Systems Handbook*, F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, Eds. Springer US, 2011, pp. 677–702.

[19] S. Amer-Yahia, S. B. Roy, A. Chawlat, G. Das, and C. Yu, "Group recommendation: Semantics and efficiency," *Proc. VLDB Endow.*, vol. 2, no. 1, pp. 754–765, Aug. 2009.