

Applying cooperative negotiation methodology to group recommendation problem

Pavel Bekkerman¹ and Sarit Kraus² and Francesco Ricci³

Abstract. In this paper we present a Group Recommender System relying on the application of cooperative negotiation methodologies. We describe a process in which automated negotiation agents, acting on behalf of human group members, participate in a cooperative negotiation, and based on individual recommendations and user preference models, generate group recommendations.

1 INTRODUCTION

In our research we study individual recommender systems integration with various cooperative and non-cooperative negotiation paradigms as a mean to introduce efficient (with emphasis on processes automation and quality) customer-customer and customer-supplier interactions in a multi-agent eCommerce environment [1].

In this paper we present an approach of applying cooperative negotiation methodology to group recommendation problem solution.

A group of users, each having an individual recommender system of his own, and based on a catalog of available products, are willing to examine possibilities and to arrive to an agreement acceptable by the whole group⁴.

Figure 1. A typical scenario

We address this typical scenario and describe a process in which automated negotiation agents, acting on behalf of human group members, participate in a cooperative negotiation, and based on individual recommendations and user preference models, generate group recommendations. This process can be further extended into a complete Group Recommender System.

2 EXISTING METHODOLOGY

We start with a short survey of the current state-of-the-art of Group Recommender Systems. While the problem of producing personalized recommendations to a single user is a field with a rich literature [2], the Group Recommender Systems is a relatively new field.

For instance, [3] specifies that as of 2001 there were no previously published studies.

The review of existing literature in Group Recommender System that we conducted has identified two main approaches to generating group recommendations from individual recommendations: *averaging* and *merging*. We decided to cover here two representing works. The work on the averaging method is represented by [4], where [3] argues for the merging method, however mentioning the pros. and cons. of the two methods.

In [4] the authors describe a Group Recommender System called Travel Decision Forum. In this system each user is able to provide its rating (discrete, 1-5) for each of the attributes of a traveling asset. Altogether, these ratings are named user-specification and are used to produce individual user recommendations. The paper further describes a method of aggregating several user-specifications into an “average” user-specification and then producing individual recommendation for this “average” user which are, finally, taken as group recommendations.

In [3] the authors present a PolyLens group recommender system, which is an extension of a MovieLens individual recommender system by the same authors. MovieLens is a collaborative filtering-based recommender, where each user is able to provide a per-movie rating (discrete, 1-10). For PolyLens, the authors elaborate individual recommendations produced for each member of the group into a single ranked list, where each movie is attributed with individual recommendation score, and resort it according to an artificially introduced social value function (merging strategy): maximize the most happy member score, maximize the least happy member score, etc. The last is the strategy actually used with justification of small group size of 2-3 people planning to watch a movie together.

An interesting point, raised by [3] is a possibility of introducing an unequal weighting of individual recommendations during the merging. For example, should we be giving a more “experienced” user more weight during the merging?

To summarize the survey, the literature and the described representing works draw a picture of a much unexplored research field, and rise many research issues that are yet to be studied.

3 PROPOSED METHODOLOGY

We proceed by presenting a Group Recommender System relying on the application of cooperative negotiation methodologies. We propose a process in which automated negotiation agents, acting on behalf of human group members, participate in a direct or mediated cooperative negotiation, and based on individual recommendations and user preference models, generate and form group recommendations.

¹ Caesarea Rothschild Institute, Department of Computer Science, University of Haifa, Haifa 31905, Israel, pavel@cri.haifa.ac.il

² BarIlan University, RamatGan 52900, Israel, sarit@cs.biu.ac.il

³ eCommerce and Tourism Research Laboratory, ITC-irst, via Solteri 38, Trento 38100, Italy, ricci@itc.it

⁴ For instance, in case of traveling information that agreement would be a set of traveling assets, and, consequently, a complete travel plan for the group.




Figure 2. A Group Recommender System

The Figure 2 draws a schema of the proposed process. In the following sections we describe in details the phases of the process in a bottom-up order.

3.1 Catalog

A catalog holds a set of available products which are being recommended and, consequently, negotiated over. This is rather a broad notion since the catalog may represent: a) a single supplier database of products; or b) a group common knowledge about products available on the market; or c) a dynamic, ever-changing, global market-place, such as the Internet.

We assume that the catalog products may be of a finite number of types, where each type defines the structure of an instance of a product in terms of its schema: attributes, their possible values, and composition.

3.2 Individual Recommendations

An Individual Recommendation System provides recommendations (ranking) of the products from the catalog. A notable constraint on the nature of the recommendation process: *the recommendations must be personalized*. A Recommender System, upon receiving user identification, must produce individual and distinct recommendations about the catalog products.

In our research, we use Trip@dvice – a personalized and self-learning recommender system. Trip@dvice is a case-based recommender for tourism assets, which utilizes twofold similarity process. Roughly: first, rank other’s cases by their similarity to the user’s previous cases; second, rank products by their similarity to products from most similar other’s cases [5].

3.3 Individual Utility

Recommendations are being invoked by negotiation agents (N.A. on Figure 2) acting on behalf of actual human users – the group

members. In addition to the recommendations (a ranked list of products), an individual utility of each product to the actual user must be evaluated. The individual utility evaluation to an actual user is the essence of individual recommendations acceptance problem. In practice, however, the human user acceptance issue in recommendation systems research is often avoided [4].

In our case, not only we must have the individual utility evaluation in order to proceed to negotiation based on this utility value, but we also need to *automate the evaluation process*. To achieve this we introduced a user preference model. A robust user model might learn from an actual user evaluation, and the model quality is only a function of a learning algorithm. Willing to demonstrate the proposed process at an early stage of our research, we took a more direct approach: modeling user preferences through a fixed multi-attribute utility function [6, 7].

Additionally, both the individual utility values and the individual recommendation values⁵ should be taken in to account in a negotiation phase of the group recommendations formation. It’s an interesting observation that taking in account only recommendations would margin to the described above existing methodology of merging, and taking in account only the individual utility would margin to a traditional cooperative negotiation.

3.4 Negotiation

A field of cooperative negotiation captures a wide range of agent interactions aiming to produce a mutually beneficial agreement(s) [1, 8]. Negotiation process formalization has many aspects [9]. Fixing these aspects in a domain of particular problem results in what is referred as a negotiation paradigm.

Several negotiation paradigms were proposed for solving the group recommendation problem. We break their description into three most influential aspects: protocol, mediation, and agreement.

1.1.1 Protocol

Negotiation protocol governs the way negotiation may develop by defining which messages that negotiation participants may send to each other, the order of message communication, etc. [9].

For the groups of two users (each represented by a negotiation agent) we studied, through simulation, several protocols: “AlternatingOffers” (direct negotiation) and “MergingRanks” (mediated negotiation).

In “AlternatingOffers” protocol each negotiation participant has an opportunity to: a) place an offer, b) accept one of the previously placed offers by the other negotiation participants. A more detailed description of the protocol and strategies for the negotiation participants is given by [7].

In “MergingRanks” protocol each negotiation participant supplies its complete ranked list to a negotiation mediator that evaluates each proposal and forms a composite (“merged”) ranked list. We provide more details on this evaluation process in the following section on negotiation mediation and the section describing the mediation strategy in the conducted experiments.

For groups of three and more users the direct communication execution and control becomes more complex, which is evident also from a real-life users’ experience. This complexity itself may bring to a failure to produce a negotiation agreement. Mediated

⁵ By the individual recommendation values we mean: whatever information the individual recommender is able to provide to justify its ranking, such as, the product scores. For example, in Trip@dvice we’re provided with similarity scores of the recommended products.

negotiation is a more obvious way to manage the negotiation and form an agreement - group recommendations.

1.1.2 Mediation

The negotiation mediator benefits from being able to simultaneously access all the proposals from all the group members. To generate an agreement, the mediator may have a “stock” of strategies to help him choosing among these proposals the one that will be offered to the group as an agreement.

Some simple strategies may be: “maximizing the utility of an average group member”, “maximizing the utility of the least happy member”, “maximizing the utility of the most happy member” [3].

In general, provided a model of group utility (social welfare), the mediator’s strategy should be: *choose among the proposal the one that maximizes the group utility.*

1.1.3 Agreement

Upon termination of the negotiation phase, an agreement is produced. With adequate choice of negotiation protocol the process also allows us to generate the second-best agreement, the third-best agreement, and so on. For example, the “MergingRanks” protocol will produce a ranked list of products, roughly,

$$\begin{aligned}
 i_1^* &\equiv \arg \max_{i \in [P]} \{u_{med}(p_i)\} \\
 i_2^* &\equiv \arg \max_{i \in [P] \setminus \{i_1^*\}} \{u_{med}(p_i)\} \\
 i_3^* &\equiv \arg \max_{i \in [P] \setminus \{i_1^*, i_2^*\}} \{u_{med}(p_i)\} \\
 &\dots \\
 i_n^* &\equiv \arg \max_{i \in [P] \setminus \{i_1^*, \dots, i_{n-1}^*\}} \{u_{med}(p_i)\}
 \end{aligned} \quad (1)$$

Where:

- P A set of products to rank
- $i \in [P]$ A unique index of the unranked product (i.e. product- i)
- $u_{med}(p_i)$ A mediator-calculated utility value for the product- i
- $i_k^* \in [P]$ An index of the k^{th} ranked product

Similarly, the “AlternatingOffers” protocol is designed to produce a complete ranked list of products [7].

3.5 Group Recommendations

The ranked list of products resulting from the agreement formation phase is the resulting group recommendations produced by the overall process (for example, in case of “MergingRanks” protocol it is given by (1)). Just as with the individual recommendations, these group recommendations are presented to the group and its members with or without a need of a human acceptance.

3.6 Group Utility (Social Welfare)

Similarly to individual recommendations, we would like to evaluate the utility of the produced recommendations to a group as a whole and, perhaps, to each individual member of the group. Unlike with individual recommendations, which are utilized as an integral part of the described process, the group recommendation

evaluation is not a part of the process, but rather a research step with a goal to estimate the quality of resulting recommendations.

4 EXPERIMENTS

We have built a complete experimental environment to demonstrate the proposed process: from the individual recommendations extraction, through negotiation, and up to generation of group recommendations. We were using Trip@dvice as an individual recommender system and performed a comparative study of several cooperative negotiation paradigms.

For the groups of two users we *simulated*, under various setups, negotiations based on “AlternatingOffers” protocol (direct negotiation), and “MergingRanks” protocol (mediated negotiation). The negotiated group recommendations were evaluated with a modeled group utility (social welfare) function, for instance: a product of both users’ utility. The same utility function was used by the negotiation mediator in “MergingRanks” protocol based on utility only:

$$\forall i \in [P]: u_{med}(p_i) \equiv u_1(p_i) \cdot u_2(p_i)$$

Where:

- $u_1(p_i)$ An individual utility value of user-1 for the product- i
- $u_2(p_i)$ An individual utility value of user-2 for the product- i

Table 1 shows the product ranking due to the individual recommendations for user-1 and user-2, followed by the product ranking due to the negotiated group recommendations produced by agents in the protocols protocol. For example, the product ranked first in individual recommendations is product-69 for user-1 and product-48 for user-2; and the products ranked first in group recommendations resulting from negotiation are: product-2 for “MergingRanks” protocol based on utility only; product-48 for “Merging-Ranks” protocol based on rank only; product-2 for “Merging-Ranks” protocol experimental (discussed below); and product-2 for “AlternatingOffers” protocol. To save space we show only the top-20 out of a total 100 ranked products.

Table 1. Negotiated group recommendations for a group of two users (product of utilities)

Rank	Product (Negotiated) Group Recommendation						
	Individual Recommendation	User-1	User-2	M.R.(u)	M.R.(r)	M.R.(exp)	A.O.
1	69	48	2	48	2	2	
2	39	42	40	40	40	100	
3	50	6	100	2	100	40	
4	85	72	60	69	60	60	
5	40	81	81	42	81	23	
6	54	20	63	6	23	13	
7	63	60	25	100	63	79	
8	2	25	23	39	94	96	
9	52	29	6	60	75	75	
10	66	2	94	81	25	94	
11	100	88	42	63	96	81	
12	27	97	75	50	29	21	
13	47	100	29	72	13	63	
14	53	94	96	85	6	71	
15	31	40	66	25	66	25	
16	67	82	13	29	42	29	
17	75	3	79	54	79	66	
18	59	24	21	66	21	15	
19	73	78	71	94	71	6	
20	23	96	38	20	53	38	

Since the mediator in “MergingRanks” protocol based on utility directly applies (2) to evaluate proposals from the agents and ranks them according to this calculated group utility value, its results are optimal, with respect to the chosen group utility function, and serve a basis for comparison with other protocols performance. On Table 1 the top-5 group recommended products according to “Merging-

Ranks” protocol based on utility only are colored, and their position is shown in the rest of the rankings.

We were also able to improve the performance in “Merging-Ranks” protocol based on rank only by modeling the *expected* functional dependency between a product utility and its rank based on the apriori knowledge of individual utility-rank distribution. The graph on Figure 3 shows a typical utility-rank dependency that we observed. For user-1 and user-2 it shows the individual utility, and for each of the protocol executions it shows the group utility value calculated using (2), as a function of recommended product rank.




Figure 3. Utility-Rank observation for a group of two users (product of utilities)

Similar outcome is presented on Table 2 where, this time, a normalized sum of both users’ utility was used as a group utility function.

Table 2. Negotiated group recommendations for a group of two users (normalized sum of utilities)

Rank	Product					
	Individual Recommendation		(Negotiated) Group Recommendation			
	User-1	User-2	M.R.(u)	M.R.(r)	M.R.(exp)	A.O.
1	69	48	2	48	2	2
2	39	42	40	69	40	100
3	50	6	100	42	100	40
4	85	72	60	39	60	60
5	40	81	81	6	81	23
6	54	20	63	50	63	79
7	63	60	6	40	25	13
8	2	25	25	72	23	96
9	52	29	42	85	94	75
10	66	2	23	81	6	94
11	100	88	94	2	75	81
12	27	97	66	60	29	63
13	47	100	29	54	42	21
14	53	94	75	20	96	71
15	31	40	96	63	66	25
16	67	82	13	100	13	29
17	75	3	79	25	79	66
18	59	24	21	29	21	15
19	73	78	88	66	53	38
20	23	96	53	52	88	6

5 DISCUSSION

In the described experiments the similarity value of each product ranked by individual recommender system was accounted as an individual utility. Doing so, we, in fact, recreated the merging method. In the further experiments we will introduce user preference models (fixed multi-attribute utility functions) to account for the ability of group members to evaluate individual recommendations prior to proceeding to the negotiation phase.

We showed that the proposed process was able to recreate the existing merging methodology. We believe that, with introduction of individual user preference model at an early stage of group recommendations formation, our group recommender will better re-

flect the preferences of the group as a whole, increasing the chances for acceptance of the produced group recommendations.

In further experiments, this time with *non-simulated individual recommendations samples*, we will concentrate our attention on a negotiation with the following information available to the agents: a) an individual recommendation-based ranking, b) an individual recommendation-based similarity value, c) a modeled individual utility, and d) an actual individual utility.

One of the interesting open questions raised by our research is: which impact has a method used to produce individual recommendations on the described process and the resulting group recommendations. For instance, which impact has utilizing case-based reasoning individual recommender, as opposed to collaborative filtering recommender?

6 SUMMARY

We presented a Group Recommender System relying on the application of cooperative negotiation methodologies. We described a process in which automated negotiation agents, acting on behalf of human group members, participate in a direct or mediated cooperative negotiation, and based on individual recommendations and user preference models, generate and form group recommendations. We described experiments demonstrating the ability of the proposed process to produce group recommendations, discussed results of these experiments, and draw a map of the further research work.

REFERENCES

- [1] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra and M. Wooldridge. Automated negotiation: prospects, methods and challenges. *Int. J. of Group Decision and Negotiation* 10 (2) 199-215. 2001.
- [2] Adomavicius, G. and Tuzhilin, A. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734-749. 2005.
- [3] O'Connor, M., Cosley, D., Konstan, J. A. and Riedl, J. PolyLens - A Recommender System for Groups of Users, In: *Proceedings of the seventh European Conference on Computer Supported Cooperative Work:199—218*. 2001.
- [4] Anthony Jameson, More Than the Sum of Its Members: Challenges for Group Recommender Systems, *Proceedings of the working conference on Advanced visual interfaces*, May 25-28, Gallipoli, Italy. 2004.
- [5] F. Ricci, A. Venturini, D. Cavada, N. Mirzadeh, D. Blaas and M. Nones. Product Recommendation with Interactive Query Management and Twofold Similarity. *Proceedings of the 5th International Conference on Case-Based Reasoning (ICCBR 2003)*. Trondheim, Norway. June 23-26. 2003.
- [6] Von Neumann, J. & Morgenstern, O. *Theory of Games and Economic Behavior* 2nd ed., Princeton University Press, Princeton, NJ. 1947.
- [7] M. Barbuceanu and W.K. Lo. A Multi-attribute Utility Theoretic Negotiation Architecture for Electronic Commerce. *Proceedings of the Fourth International Conference on Autonomous Agents*, Barcelona, Spain. 2000.
- [8] S. Kraus. *Strategic Negotiation in Multi-Agent Environments*. MIT Press, Cambridge, USA. 2001.
- [9] C. Bartolini and C. Preist. *A Framework for Automated Negotiation*. HP Labs, Technical Report. 2001.