Preference Networks and Non-Linear Preferences in Group Recommendations

Amra Delic
TU Wien, Faculty of Informatics
Vienna, Austria
amra.delic@tuwien.ac.at

Francesco Ricci
Free University of Bozen-Bolzano
Faculty of Computer Science
Bozen-Bolzano, Italy
Francesco.Ricci@unibz.it

Julia Neidhardt
TU Wien, Faculty of Informatics
Vienna, Austria
julia.neidhardt@tuwien.ac.at

ABSTRACT

Group recommender systems generate recommendations for a group by aggregating individual members' preferences and finding items that are liked by most of the members. In this paper we introduce a new approach to preference aggregation and group choice prediction that is based on a new form of weighting individuals' preferences. The approach is based on network science, and, in particular, it relies on the computation of node centrality scores in preferences similarity networks of groups. We also motivate and introduce a non-linear (exponential) remapping of the individuals' preferences. Based on offline experiments we demonstrate: 1) non-linear remapping of preferences is useful to better predict group choices and generate recommendations; and 2) our weighted approach predicts the actual group choices more accurately than current state-of-the-art methods for group recommendations.

CCS CONCEPTS

• Information systems → Recommender systems; **• Human-centered computing → User studies**; *User models*; *Social network analysis.*

KEYWORDS

Preference Aggregation, Network Science, Centrality

ACM Reference Format:
Amra Delic, Francesco Ricci, and Julia Neidhardt. 2019. Preference Networks and Non-Linear Preferences in Group Recommendations. In *IEEE/WIC/ACM International Conference on Web Intelligence (WI '19), October 14–17, 2019, Thessaloniki, Greece.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3350546.3352556

1 INTRODUCTION

Research in group recommender systems (GRSs) usually deals with the question of how to correctly aggregate group members' individual preferences in order to deliver recommendations that the group will actually chose and are fair. As Arrow's theorem indicates and the GRSs research has shown, no single, optimal approach to aggregate individuals' preferences exists [20]. Moreover, the classical preference aggregation strategies are not uniformly adopted by groups, i.e., different groups adopt different approaches to make their decisions. Hence, having the goal to predict or reconstruct the actual outcome of a group decision process, new approaches are required in order to generate more precise recommendations. In this work, we adapt a general strategy, i.e., the classical average aggregation method, by weighting differently the importance of the group members when their preferences are aggregated. We introduce two novel types of adaptations of general preference aggregation strategies: a) a novel method to weight individuals' preferences in a group preference model before they are aggregated by any standard method; b) a non-linear transformation of the original individual preferences of the group members before they are aggregated in a standard strategy.

A considerable amount of work has been done in the field of psychology to better understand group behavior in decision-making processes. It was shown that these processes are influenced by different types of connections (relationships) that emerge between pairs of group members [10], and therefore by different types of group networks. In this paper, we analyse preference similarity networks, which evolve when individuals with various preferences are placed in a group. We define network connections (i.e., edges) on the basis of group members' pairwise preference agreement derived from the Spearman Footrule Distance. As in any other network, the "network centrality" [26] of an actor plays an important role for the development and the outcome of various processes within that network. Similarly, we hypothesize that, firstly, the centrality of a user in the group preference network should have an impact on the individual's weight in the decision-making process. Secondly, we propose a method to weight group members' individual preferences based on their centrality in that network in order to improve the prediction of group choices.

Furthermore, we challenge the assumption of linearly discounting lower ranked choice preferences in the preference aggregation strategy. In fact, Masthoff [19] has observed that the prediction of a group choice is improved when the linear rating scale is transformed into a quadratic one. Moreover, in many cases humans actually react linearly to an exponential growth of the stimulus [25]. This is the reason why log scales are typically used (e.g., decibel or lumen). Thus, we assume that a non-linear (exponential) remapping of the individuals' preferences can improve the prediction of the group choices.

In conclusion, here we make two hypotheses and we test them by using two data sets describing the preferences and the choices of 282 participants in 79 real groups while deciding on a travel

WI '19, October 14–17, 2019, Thessaloniki, Greece

ACM ISBN 978-1-4503-6934-3/19/10…$15.00
https://doi.org/10.1145/3350546.3352556

destination to visit together: ***H1*** Transforming the scale of group members' individual preferences from a linear to a non-linear one, improves the prediction of group choices. ***H2*** Using group members' centrality in a group preference network in order to weight their importance in a group preference model, improves the prediction of the group choices.

The results of experiments show that our first hypothesis is partially confirmed, and that our second hypothesis is strongly supported by the obtained results.

The rest of the paper is organized as follows: in Section 2 we provide an overview of the related work. In Section 3 the data collection process is explained; and in Section 4 the methods to predict the group choice are introduced. In Section 5 we explain how we evaluate the performance. In Section 6 the results are presented, and in Section 7 we discuss these results and present our conclusions.

2 RELATED WORK

Many aggregation strategies have been proposed in the literature [20], mostly motivated by the Social Choice Theory, including *Borda Count* that aggregates individuals' ranked lists into a group list and *(Weighted) Average* that averages individuals' item-ratings. In certain cases preferences of some group members are considered as more important, then a weighted average is used. *Multiplicative* multiplies individuals' item-ratings, and *Least Misery* assumes that a group rating is the minimum of individuals' ratings. In [9, 19], the authors have evaluated whether or not people, when generating group recommendations and making decisions, would follow a clear strategy. The results have shown that this was in fact the case, but there was no single, dominating strategy employed.

Another line of research in GRSs deals with role-based models in which members are weighted based on the group role they have. For instance, roles are derived from members' activity in the system [1, 4], domain expertise [11], strength of personality [23], social centrality and similarity [2], or based on the previous group choices [24]. Our approach is also role-based, where the role / weight of a group member is defined by her "preference centrality / status" in the group. The idea emerged from psychology literature, where Kameda et al. [15] have shown that cognitive centrality (i.e., a member's status in the group determined by the level of information and knowledge that a member shares with the rest of the group) plays a significant role when identifying sources of influence and predicting outcomes of the group decision-making process. Hereby we use the notion of "preference status" in GRSs. Practically, "preference status" is extracted from the group preference network as the measure of centrality. Then, it is used in a weighted aggregation strategy to deliver group recommendations (i.e., predict actual group choices).

An approach similar to ours has been employed, in [18], but for addressing a different task: a group formation problem based on the pairwise similarity of users' ratings.

3 GROUP DATA

The observational data used in this paper was collected in a user study focused on the travel and tourism domain. The first implementation of the study took place at TU Wien, Uni Klagenfurt, TU Delft and Uni Leiden in 2015 and 2016, and the second one at TU Wien and Uni Sarajevo in 2017. Both implementations followed the same three-phases structure [6–8] that we present in the following.

In the first study phase individual preferences were collected with a so-called pre-questionnaire, i.e., participants' explicit ratings or rankings of ten pre-selected destinations. In fact, the two implementations differed in the pre-selected destinations, and in the way participants expressed their preferences about them. In the first implementation the destinations were ten large European cities and the participants *ranked* them. Conversely, in the second implementation, the destinations were chosen to fit the general preferences of certain traveller types identified in the tourism literature [12, 13, 21, 22, 27] and the participants *rated* them on a ten-point scale (1 - not attractive at all, 10 - highly attractive). In the second phase of the study, the participants formed groups, and selected, from the pre-defined set, two top destinations that they, as a group, would like to visit together. Finally, in the third phase, the participants filled in a post-questionnaire, where they reported their first and second group choice, and answered a number of questions including individual choice satisfaction and perceived difficulty of the process.

The implementations resulted in two data sets: DSI (200 participants in 55 groups), and a smaller set DSII (82 participants in 24 groups). The groups had a minimum of two and a maximum of five members. In this work, we focus entirely on participants' explicitly formulated individual preferences and their group choices.

4 GROUP CHOICE PREDICTION METHODS

Group choice prediction methods are algorithms that aggregate the individual preferences of the group members, and generate a ranking or rating for the options. These ratings or rankings are used to *predict* the actual choice of the group (or to generate group recommendations).

To test our **first hypothesis**, we transform the group members' individual preferences from a linear scale into a non-linear one. We use an exponential function, as shown in equation 1:

$$new\_rating = a^{old\_rating} \tag{1}$$

where the parameter a is set to 1.5, in order to simulate a mild exponential growth. In this paper we do not analyse the effect of different values for the base a, we leave this for future work.

To test our **second hypothesis** we introduce preference networks and accordingly propose a method to calculate group members' weights based on their centrality in those networks.

Preference Network. For each group, a separate network is constructed. The nodes represent group members, and the edges represent relationships between them. These relationships are quantified with node-to-node similarities, which were obtained from a metric similar to the Spearman's Footrule Distance function [3]. In fact, we use the Full Choice-set Distance measure (*FullDist*). It considers members' preferences for the full set of options (*ChoiceSet*), to compute the distance between two group members u and v. It gives an undirected preferences relationship between pairs of group members:

$$FullDist(u, v) = \sum_{i \in ChoiceSet} |score_u(i) - score_v(i)| \tag{2}$$

where $score_u(i)$ is: in DSI (and DSII respectively) the position (rating) of option i in the personal ranking (ratings) of the ten considered options of a group member u. Hence this score is in both cases (rankings and ratings) an integer between 1 and 10.

To represent the relationship between two group members u and v with a similarity rather than a distance, we simply subtract the distance score between u and v from the overall maximum distance score obtained.

Network Analysis. Representing the preferences of group members and their relationships in such a manner allows us to apply network analysis methods at the group level, and to define the group members centrality scores. Semantically, a group member with a high centrality score, in such a network, is a person that shares a great deal of preferences with the rest of the group. This group member could then be seen as more important or prominent since this member is "positioned" at the intersection of all the group members' preferences. We use the degree centrality of a node, which is the sum of all its links' weights. Although there are other centrality measures [26], this is the most obvious choice as we consider rather small, fully connected, undirected, and weighted network.

Then, the weight w_u of a group member u is calculated by normalizing the centrality score $CS(u)$ of user u in its group preference network G:

$$w_u = \begin{cases} 1 & \text{if } \max_G(CS) = \min_G(CS) \\ \frac{CS(u) - \min_G(CS)}{\max_G(CS) - \min_G(CS)} + \beta, & \text{otherwise} \end{cases} \tag{3}$$

where $\max_G(CS)$ and $\min_G(CS)$ are, respectively, the maximum and the minimum node centrality scores within the graph G. The parameter β controls the relative differences of the group members' weights (larger values will make all the weights ratios closer to 1).

Baseline Prediction Methods. Since we hypothesize that the two proposed approaches can improve the prediction of group choices, we first select the baseline methods that we will use for comparison. Based on the individual preferences (ratings or rankings) of the group members, these strategies generate either a ranking or a scoring for the available options. Several aggregation strategies were introduced in Section 2, and according to the related literature, we will consider average, multiplicative, least misery and Borda count as the baseline methods since these are the methods that were used most often.

Apart from Borda that operates on the rankings, these strategies require, in order to compute the group score, the group members' ratings as the input. As previously mentioned, the two data sets, DSI and DSII, were different in terms of group members individual preferences, i.e., rankings in the first vs. ratings in the second. Hence, to apply the Borda count strategy on the DSII data set, we simply mapped the ratings into partial rankings of the options. Furthermore, to apply the average, multiplicative and the least misery approach on the DSI data set, we required for each user the individual ratings of the ten destinations. Hence, in this case the rankings were mapped to ratings, by applying the two steps procedure: 1) Given the individual partial rankings calculated for DSII (as previously explained), we calculated for each ranking position the conditional probabilities for each rating that it will be assigned to the rank at hand; 2) For each group member from DSI and for each ranking that she provided, a rating was sampled according to the previously calculated conditional probabilities.

Finally, with the goal to test our research hypotheses, we introduce two categories of preference aggregation strategies, that are compared with the baseline methods:

(1) The non-linear category including: non-linear Borda (*"Nl-Borda"*), which uses non-linear rankings of the group members; non-linear average (*"Nl-Avg"*), and non-linear multiplicative strategy (*"Nl-Mult"*), which use the non-linear ratings.
(2) The weighted category including: weighted average (*"W-Avg"*); weighted Borda count (*"W-Borda"*); and weighted multiplicative (*"W-Mult"*) (group members' weights are calculated according to the equation 3).

5 EVALUATION STRATEGY

To asses the predictive performance of the aggregation strategies, i.e., the baseline and the newly proposed ones, we used precision of the recommendations at top-one ($P@1$) and top-two ($P@2$):

$$P@k = \frac{\#CorrectTopkRecommendations}{\#TopkRecommendations} \tag{4}$$

The numerator captures the number of group choices that are correctly predicted by the aggregation strategy. The denominator is the number of top-k recommendations generated by an aggregation strategy (i.e., when multiple items obtain the same score they will share the rank in a recommendation list, but the precision score will decreased accordingly).

As the weighted aggregation strategies depend on the β parameter (in equation 3), we searched for the β value which maximizes the estimated precision of the prediction strategy. To this end, we have split the data into a training set (randomly selected 50% of groups from the combined data set, i.e., DSI and DSII merged), and a test set (the remaining 50% of the groups). In the training set, the β value, for which the considered weighted aggregation strategy reached its maximum precision score (i.e., the mean value of $P@1$ and $P@2$) was selected. The test set was used to compare all the proposed methods with the baselines. The procedure of randomly splitting the data was repeated all together ten times and the average performance metrics were computed. In each repetition the optimal β parameter was selected based on the results obtained for the training set.

6 RESULTS

6.1 Non-linear Scaling of Preferences

In the first hypothesis we conjecture that the prediction of the group choices can be improved if group members' individual preferences are transformed with an exponential function, hence are mapped from a linear scale to a non-linear one. Table 1 shows the precision of the considered strategies, when predicting the two choices of the group with the baseline and non-linear methods. Here, the best scores are highlighted, where the higher is the better. As previously explained, the performance metrics are computed on a test set formed by a random sample of 50% of the groups, after the weighted methods were trained on the other 50% of the groups for selecting the optimal β value. Since this procedure was repeated ten times, we report the average precision scores over these ten repetitions.

Table 1: Non-linear Scaling - Average P@1 and P@2

Aggregation strategy	P@1	P@2
Average	0.458	0.422
Nl-Average	0.518	0.450
Borda	0.550	**0.449**
Nl-Borda	**0.608**	0.445
Multiplicative	0.445	0.426
Nl-Multiplicative	0.458	0.422
Least misery	0.414	0.358

The $P@1$ scores of average and Borda strategies are improved by the exponential (non-linear) mapping of the group members' individual preferences. The improvement of the non-linear Borda in comparison to the baseline Borda is over 5%, while the improvement of the non-linear average strategy is 6%. The non-linear multiplicative also performs better better than its baseline, but only slightly (1.3%). The least misery strategy will not be considered anymore, since in this data set it performs worse than all the other methods. For $P@2$, the proposed non-linear mapping improves only the average strategy, with a slight improvement of 2.8%. The non-linear Borda and multiplicative perform almost the same as their respective baselines.

To summarize, the non-linear Borda performs particularly better than the baseline Borda with respect to the $P@1$ metric, but this does not hold for $P@2$. However, the non-linear average strategy obtains notably better results in both cases. As for the non-linear multiplicative strategy, no clear improvement is achieved for both precision metrics.

6.2 Centrality-based Weighting

In the second hypothesis we conjecture that weighting group members according to their centrality in the group preference network can improve the precision of an aggregation strategy prediction of the group choices in comparison to the corresponding baseline strategy.

Table 2: Centrality-based Weighting - Average P@1 and P@2

Aggregation strategy	P@1	P@2
Average	0.458	0.422
W-Average	0.523	0.438
Borda	0.550	0.449
W-Borda	**0.603**	**0.455**
Multiplicative	0.445	0.426
W-Multiplicative	0.525	0.454

For P@1 the weighted average performs better than the baseline average (see table 2). The weighted Borda strategy achieves the second best $P@1$ result of all the aggregation strategies (only slightly weaker than Nl-Borda). The weighted multiplicative also preforms better than its baseline, even by 8%. For $P@2$, the improvement of the weighted methods is more noticeable. The weighted average is again better than the baseline average (table 2). The weighted Borda is the best performing strategy in this case, and the weighted multiplicative is the second best strategy with almost 3% improvement in comparison to its baseline.

To summarize, the weighted aggregation strategies perform consistently better than their respective baselines, achieving better $P@1$ and $P@2$ scores, on average.

7 DISCUSSION AND CONCLUSION

In this section we discuss the implications of our results. To test the first hypothesis we evaluated the performance of certain preference aggregation strategies when the group members' preferences are transformed from a linear scale to a non-linear one using an exponential function. The hypothesis is motivated by the fact that, in many situations, humans react linearly (with the given rating or ranking) to an exponential growth of the stimulus (attractiveness), as described in section 1. In other words, we assume that for a certain option to be one level better than another option it must have a score that is larger than the other by a multiplicative factor. The results we obtained partially support this hypothesis. The P@1 metric has clearly improved for all the evaluated strategies when considering a non-linear scaling of the preferences, while the P@2 metric has only improved for the non-linear average strategy. Certainly, the performance of the non-linear strategies is influenced by the selection of the parameter a (i.e., the base of the exponential function), which we set to 1.5 in the current analysis. For the future work, we plan to evaluate the effect of the a parameter, and optimize it.

In the second hypothesis, we conjecture that representing groups in a network structure and using the degree centrality of the nodes to weight the respective members in a preference aggregation strategy will improve the precision of the group choices prediction. Such an approach is used in group dynamics theory [10, 15–17], and has recently been evaluated for the group formation problem [18]. However, to our best knowledge, such an approach has not yet been attempted in GRSs. According to our results, this *hypothesis is supported and improvements can be achieved when using the weighted strategies instead of their respective baselines*. This is true for both precision metrics (P@1 and P@2).

The performance of the weighted strategies is highly dependent on the proper selection of the β parameter. In our approach, for each training set, the best performing β was selected. Clearly, the larger the training set, the greater the confidence to select β. However, our data set comprises only 79 instances, 50-50 split of the data results in 39/40 training instances only. We could have used a different splitting ratio (e.g., 70-30), but in that case the number of groups on which we test the performance of the aggregation strategies becomes severely limited (e.g., containing only 23/24 groups).

We are aware that the size of the data set is a limitation of this work, and in the future, we are planning to test our hypotheses on other data sets, such as those used in [5, 14]. Nevertheless, showing that our methods perform on average better than the baselines, strongly support the validity of our second hypothesis.

REFERENCES

[1] I. Ali and S. W. Kim. 2015. Group recommendations: approaches and evaluation. In *Proceedings of the 9th International Conference on Ubiquitous Information Management and Communication*. ACM, 105.

[2] I. Alina Christensen and S. Schiaffino. 2014. Social influence in group recommender systems. *Online Information Review* 38, 4 (2014), 524–542.

[3] L. Baltrunas, T. Makcinskas, and F. Ricci. 2010. Group recommendations with rank aggregation and collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems, RecSys'10*. ACM, Barcelona, Spain, 119–126.

[4] S. Berkovsky and J. Freyne. 2010. Group-based recipe recommendations: analysis of data aggreagation strategies. In *Proceedings of the 4th ACM conference on Recommender systems*. 111–118.

[5] D. Cao, X. He, L. Miao, Y. An, C. Yang, and R. Hong. 2018. Attentive group recommendation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. ACM, 645–654.

[6] A. Delic, J. Masthoff, J. Neidhardt, and H. Werthner. 2018. How to Use Social Relationships in Group Recommenders: Empirical Evidence. In *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization*. ACM, Singapore, Singapore, 121–129.

[7] A. Delic, J. Neidhardt, T. N. Nguyen, and F. Ricci. 2018. An observational user study for group recommender systems in the tourism domain. *Information Technology & Tourism* 19, 1-4 (2018), 87–116.

[8] A. Delic, J. Neidhardt, T. N. Nguyen, F. Ricci, L. Rook, H. Werthner, and M. Zanker. 2016. Observing Group Decision Making Processes. In *Proceedings of the tenth ACM conference on Recommender systems, RecSys'16*.

[9] A. Delic, J. Neidhardt, and H. Werthner. 2018. Group Decision Making and Group Recommendations. In *2018 IEEE 20th Conference on Business Informatics (CBI)*, Vol. 1. IEEE, Vienna, Austria, 79–88.

[10] D. R. Forsyth. 2014. *Group Dynamics* (6th ed.). Wadsworth Cengage Learning.

[11] M. Gartrell, X. Xing, Q. Lv, A. Beach, R. Han, S. Mishra, and K. Seada. 2010. Enhancing group recommendation by incorporating social relationship interactions. In *Proceedings of the 16th ACM international conference on Supporting group work*. ACM, 97–106.

[12] H. Gibson and A. Yiannakis. 2002. Tourist roles: Needs and the lifecourse. *Annals of tourism research* 29(2) (2002), 358–383.

[13] U. Gretzel, N. Mitsche, Y.-H. Hwang, and D. R. Fesenmaier. 2004. Tell me who you are and I will tell you where to go: Use of travel personalities in destination recommendation systems. *Information Technology & Tourism* 7, 1 (2004), 3–12.

[14] F. Hu, X. Huang, X. Gao, and G. Chen. 2019. AGREE: Attention-Based Tour Group Recommendation with Multi-modal Data. In *International Conference on Database Systems for Advanced Applications*. Springer, 314–318.

[15] T. Kameda, Y. Ohtsubo, and M. Takezawa. 1997. Centrality in sociocognitive networks and social influence: An illustration in a group decision-making context. *Journal of personality and social psychology* 73, 2 (1997), 296.

[16] A. W. Kruglanski and D. M. Mackie. 1990. Majority and minority influence: A judgmental process analysis. *European review of social psychology* 1, 1 (1990), 229–261.

[17] B. Latane. 1996. Strength from weakness: The fate of opinion minorities in spatially distributed groups. *Understanding group behavior: Consensual action by small groups* 1 (1996), 193–219.

[18] H. Mahyar, E. Ghalebi K., S. M. Morshedi, S. Khalili, R. Grosu, and A. Movaghar. 2017. Centrality-based group formation in group recommender systems. In *Proceedings of the 26th International Conference on World Wide Web Companion*. International World Wide Web Conferences Steering Committee, 1187–1196.

[19] J. Masthoff. 2004. Group modeling: Selecting a sequence of television items to suit a group of viewers. In *Personalized digital television*. Springer, 93–141.

[20] J. Masthoff. 2015. Group recommender systems: aggregation, satisfaction and group attributes. In *Recommender Systems Handbook* (2nd ed.), F. Ricci, L. Rokach, and B. Shapira (Eds.). Springer, 743–776.

[21] G. Moscardo, A. M. Morrison, P. L. Pearce, C.-T. Lang, and J. T. O'Leary. 1996. Understanding vacation destination choice through travel motivation and activities. *Journal of Vacation Marketing* 2, 2 (1996), 109–122.

[22] J. Neidhardt, R. Schuster, L. Seyfang, and H. Werthner. 2014. Eliciting the users' unknown preferences. In *Proceedings of the 8th ACM Conference on Recommender systems*. ACM, 2645767, 309–312. https://doi.org/10.1145/2645710.2645767

[23] L. Quijano-Sanchez, J. A. Recio-Garcia, B. Diaz-Agudo, and G. Jimenez-Diaz. 2013. Social factors in group recommender systems. *ACM Transactions on Intelligent Systems and Technology (TIST)* 4, 1 (2013), 8.

[24] E. Quintarelli, E. Rabosio, and L. Tanca. 2016. Recommending New Items to Ephemeral Groups Using Contextual User Influence. In *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, 285–292.

[25] Stanley Smith Stevens. 2017. *Psychophysics: Introduction to its perceptual, neural and social prospects*. Routledge.

[26] S. Wasserman and K. Faust. 1994. *Social network analysis: Methods and applications*. Vol. 8. Cambridge university press.

[27] A. Yiannakis and H. Gibson. 1992. Roles tourists play. *Annals of tourism Research* 19, 2 (1992), 287–303.