

What is the role of Context in Fair Group Recommendations?

Sara Migliorini, Elisa Quintarelli, Damiano Carra, and Alberto Belussi

Department of Computer Science, University of Verona, Italy
{name.surname}@univr.it

Abstract. We investigate the role played by the context, i.e. the situation the group is currently experiencing, in the design of a system that recommends sequences of activities as a multi-objective optimization problem, where the satisfaction of the group and the available time interval are two of the functions to be optimized. In particular, we highlight that the dynamic evolution of the group can be the key contextual feature that has to be considered to produce fair suggestions.

Keywords: Sequence recommendations · Context · Fairness · Groups

1 Introduction

Recommender systems (see the surveys [2] and [12]) have quickly become a fundamental technology, employed by many important companies such as Netflix, Amazon and Google, since they suggest items, chosen from big datasets, that match the tastes of users. Most of these systems mainly focus on personal preferences without considering the balance with fairness, which abstractly means to not discriminate against individuals or groups, when providing suggestions either to single users or groups. In all the scenarios where activities are inherently social (e.g. going to the cinema, eating out or visiting a city) recommendation systems provide suggestions for **groups** of users, which in the literature are classified either as persistent or ephemeral [9,4,6]. Persistent groups are those where members have a history of activities together, ephemeral groups, on the other hand, may be constituted by people who are together for the first time: consequently, the group preferences have to be computed on the basis of similar persistent groups, where group similarity is evaluated by means of common features defining the context of the groups [5].

The online recommendation of the next activity, either for single users, or for groups, has been extensively studied [9,2,14]: we refer to this type of recommendation as *myopic*, since it focuses on one activity at a time. The authors in [14] propose an algorithm that maximizes the satisfaction of each group member while minimizing the unfairness between them. However such an approach does not consider that further recommendations could produce a better satisfaction of the group preferences by extending the evaluation to an interval of time (optimization problem).

The recommendation of a **sequence** of activities has received little attention. For instance, the literature considers some scenarios such as the set of points of interest

for tourist [13], or the list of songs to listen [11], in which planning ahead the recommendation of a whole sequence – given some constraints, such as the available time – provides more flexibility. Moreover, it gives the chance to find solutions that a myopic recommendation may not be able to find. Nevertheless, most of the works for sequence recommendations focus on a single user and do not consider groups (see [11] for a recent survey). There are only few papers that study such a scenario [13,10,3]. In [13], the authors provide a system that suggests the path to follow within a museum by a group of visitors. The authors of [10] propose a method for suggesting a sequence of songs to a group of listeners trying to balance the users satisfaction levels. Herzog [3] considers the Tourist Trip Design Problem (sequence of points of interests) for a group of users. All the above works share a common limitations: they consider a single utility function for each user. In our work, instead, we consider the case where the choice is driven by multiple criteria, and formalize the problem as a multi-objective optimization problem.

In this paper we consider the problem of recommending a sequence of activities to groups of users, either persistent or ephemeral. The resulting sequence could reduce the thinking time required by a myopic recommendation system – e.g., channel surfing that users typically perform to find something interesting to watch next, or discussions during a trip to decide the next thing to visit. Recommending sequences can be useful whenever the group of users has a limited time interval to spend together, since suggesting a sequence decreases the time wasted in selecting the best next activity. In addition, we consider the role played by the **context**, i.e. the situation, the group is currently experiencing. It has been recognized that the notion of context influences the user preferences [1], which may vary with respect to the group current context. For instance, the same person, may prefer to see different film genres depending on the people she/he is with (e.g., family, or friends), or depending on the day (weekdays, weekend), or on period of the year (e.g., Christmas time).

Our approach proposes to model the generation of recommendation sequences as a multi-objective optimization problem, where the satisfaction of the group is one of the functions to be optimized in order to provide fair recommendations and maintain the group as cohesive as possible; in this case, fairness is understood in terms of how well our approach respects the individual preferences of the group members. We propose to solve such problem by using the Multi-Objective Simulated Annealing (MOSA) optimization heuristic [7]. More specifically, we extend such approach with the notion of context*. The novelty with respect to the state of the art (see [11]) is that time (a contextual feature) influences the group composition, thus the current context. We envision that, given the sequence recommendation problem, time has a great impact in the evolution of groups and has a central role in the optimization of user/group preferences in order to achieve fairness.

2 Context and Motivation

The proposed approach is general enough to be applied to any activity; however, due to the availability of a dataset related to the TV domain, we have considered such an entertainment scenario and its issues as our motivating example.

*The extended version of the paper has been published in [8].

Sequences are produced for groups evolving along with time. We are interested in generating sequences of recommendations for people spending a time interval together. The composition of the group, which can change over time, is a feature that influences the recommendations, since movies to be suggested to adults are not always proper for kids. We consider as contextual features the group composition (*adults with kids*, *adults*, *teens*, or *kids*) and the temporal information (*daytime* or *night*, and the day of the week), but the proposal can be easily extended by including other relevant features to describe the contexts users may experience in a given scenario. We say that a group has changed whenever its current composition changes, i.e. from *adults with kids* the group has evolved into *adults* (no matter the number of members): our recommendations are contextual, thus, a change in context determines a different policy in providing recommendations to the group in its current state. Note that the group evolution is natural along with time, indeed, a group can dynamically change its composition; thus, a contextual feature (in this case the time) can have an impact on another feature (the group composition).

Problem to solve: To optimize recommendation sequences for a group in a time interval, it is necessary to investigate how time influences the group composition.

Constraints and objective functions are considered in exploring the problem. MOSA is a multi-objective optimization technique that can take into account various constraints to be satisfied and functions to be optimized: different alternatives can be considered in aggregating individual preferences to obtain the group recommendation. In our scenario, we will consider as constraints the maximum interval of time the group can spend together (T_{\max}) and the maximum available budget (b_{\max}). As functions to optimize we will consider: i) the minimization of the empty slot between two following TV programs; ii) the maximization of the portion of T_{\max} covered by recommendations; iii) the minimization of the number of programs the group components has already watched in the past; iv) the maximization of the group satisfaction. Notice that the latter is influenced not only by the context, but also by the possible group evolution.

Problem to solve: Find the set of functions that can best describe our scenario.

Customized sequences of recommendations for groups. We analyze different possibilities to aggregate individual preferences to generate group preferences. To compute the preferences of a group for a given program, we use the classical aggregation functions, i.e. the Average Preference or the Least-Misery Preference. However, the analysis of past evolutions of groups can allow us to predict the evolution of the current group over time, e.g. it is frequent that after 9 p.m the group changes from *adults with kids* to *adults*; in this case we try to maximize the satisfaction of kids inside the group in the early phase, since we know that with high probability the group will evolve into another one with a different composition and we want to provide fair recommendations.

Problem to solve: Consider the impact of the group evolution to provide fair recommendations and prevent voluntary desertion due to unfair suggestions.

3 A Recommendation System for Groups

Fig. 1 shows an overall picture of our approach (in the blue rectangles the steps performed at run-time): we collect the description of items in the considered scenario, e.g.

different kind of entertainments, and the logs of past accesses both for users and groups. When a group is ephemeral, its preferences will be inferred starting from the components' preferences or by using the preferences of a similar (w.r.t. the context) group. The context is related to the group composition (*adults with kids, adults, teens, or kids*)

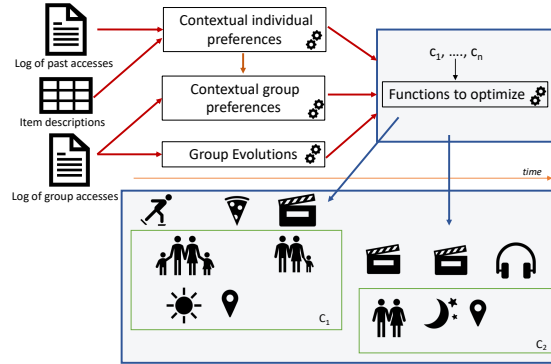


Fig. 1: Our approach: a high level vision

and temporal information (*daytime* or *night*, and the day of the week), but it may include other features. Given a group in a certain context (the temporal information can be sensed, while the group composition has to be declared) and the slot of time to spend together, our framework analyzes the preferences computed off-line and determines a sequence of suggested entertainments. These sequences are affected also by the possible group evolution computed using the past known history. In particular, if the current group is a persistent one, its past history can be used to predict its evolution with a high degree of confidence (the relative frequencies of its previous evolutions are computed, choosing the one with the higher frequency); otherwise, we can use the past history of similar groups in order to make such prediction also for an ephemeral group.

4 Evaluation

Available dataset. We consider a dataset in the TV domain, with users watching TV programs possibly in groups and in different contexts. Contexts are identified by temporal information (*daytime* or *night*, and the day of the week) and the type of groups, i.e. the number and age of members. The dataset contains TV viewings related to almost 8,000 users and 119 channels and a log containing both individual and group viewings. The log contains approximately 5 million entries, among which we retained just the synthonizations longer than three minutes. Almost 1.5 million viewings involve more than one person together. Each log row specifies the identifiers of both the user and the program he/she watched, along with start time and end time.

Dataset analysis. From the analysis of the dataset we found out that: i) the average number of programs each user watches sequentially in a day is about 3; ii) in 80% of cases where there is a sequence of program viewings for a user in a day, the number of short viewings is greater than the number of long (full) viewings and in the 46% of such

cases the number of short viewings is more than double than the number of full viewings; iv) in 32% of cases where there is a sequence of viewings for the same user in a specific day, there are at least two long viewings with short viewings in the middle. The last analysis suggested us the importance of providing sequences of recommendations also in the TV domain, since it could reduce the channel surfing activity.

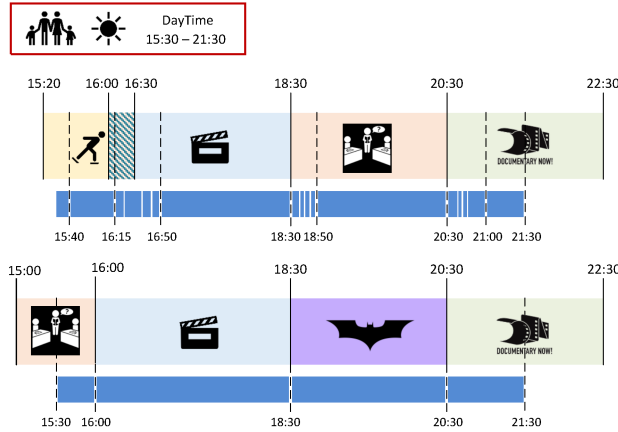


Fig. 2: Comparison of two sequences: one original and one produced by EMOSA

Fig. 2 compares a sequence contained in the original dataset and a corresponding suggestion produced by our EMOSA algorithm. Both sequences take as initial context a group of two adults and two children, in a day time period starting at 15:30 and ending at 21:30. For each sequence in Fig. 2, we report both the involved tv shows with their genre (represented by an icon) and their duration (see the timestamps above and the straight vertical lines), and the effective views performed by the group with its duration (see the timestamps below and dashed lines). In particular, the blue segments denote the duration of each single viewing, as it is noticeable in the original sequence, the group performs some short viewings between one tv show and the other (i.e., surfing activity, searching for the next TV show to view). Moreover, in the original sequence the first two chosen shows (the sport event and the film) overlap (see the mixed colours between 16:00 to 16:30). The original sequence is composed of 4 TV shows overall: the first one is a sporting event the group starts to view after 10 minutes of channel surfing and stops to view before its end; conversely, the members in the group start to watch the following film after it has already started, but they watch until the end. From this behaviour, we can conclude that the group would have seen the film from the beginning if they had known about it; on the contrary, they had less interest about the TV show, since they watched it only for a small fraction.

The sequence produced by the EMOSA algorithm is reported in the line below: it clearly removes all short views between the long ones; moreover, no loss is registered in the view of the two central TV shows. In particular, the film contained in the original sequence can now be completely watched. Moreover, the prediction about the potential group evolution (namely the fact the children will leave the group after the 21:30), led to favour their preferences and this results into the inclusion of a cartoon after the

film, which also can be entirely watched by the group. Additionally, considering this possible group evolution, the documentary is left at the end of the required period, since it is mostly preferred by the adults who are more likely to remain in the future group (eventually continuing to watch the show until its end).

5 Conclusion

We have presented a system based on a multi-objective optimization algorithm that recommends sequence of activities to dynamic group of users considering their contextual information; it is able to propose to a group of users a sequence of entertainments that improves the objective functions w.r.t. the group choices registered in the logs. The dynamic group hypothesis allows us to further improve the obtained recommendations and to balance preferences and fairness. Future work includes the application of the solution to other application domains and the execution of additional experiments by collecting also the opinion of the users about the provided recommendations.

References

1. Adomavicius, G., Sankaranarayanan, R., Sen, S., Tuzhilin, A.: Incorporating contextual information in recommender systems using a multidimensional approach. *ACM Transactions on Information Systems* **23**(1), 103–145 (2005)
2. Adomavicius, G., 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. Herzog, D.: Recommending a sequence of points of interest to a group of users in a mobile context. In: *Proceedings of ACM RecSys'17*. pp. 402–406 (2017)
4. Jameson, A., Smyth, B.: Recommendation to groups. In: *The Adaptive Web: Methods and Strategies of Web Personalization*, pp. 596–627 (2007)
5. Masthoff, J.: Group modeling: Selecting a sequence of television items to suit a group of viewers. In: *Personalized digital television*, pp. 93–141 (2004)
6. Masthoff, J.: Group recommender systems: Aggregation, satisfaction and group attributes. In: *Recommender Systems Handbook*, pp. 743–776 (2015)
7. Migliorini, S., Carra, D., Belussi, A.: Adaptive Trip Recommendation System: Balancing Travelers among POIs with MapReduce. In: *IEEE BigData Congress*. pp. 255–259 (2018)
8. Migliorini, S., Quintarelli, E., Carra, D., Belussi, A.: Sequences of Recommendations for Dynamic Groups: What is the Role of Context? *IEEE BigData Congress* (2019)
9. O'Connor, M., Cosley, D., Konstan, J.A., Riedl, J.: PolyLens: A recommender system for groups of user. In: *ECSCW 2001*. pp. 199–218 (2001)
10. Piliponyte, A., Ricci, F., Koschwitz, J.: Sequential music recommendations for groups by balancing user satisfaction. In: *UMAP workshops* (2013)
11. Quadrona, M., Cremonesi, P., Jannach, D.: Sequence-Aware Recommender Systems. *ACM Comput. Surv.* **51**(4), 66:1–66:36 (2018)
12. Ricci, F., Rokach, L., Shapira, B.: Recommender systems: introduction and challenges. In: *Recommender Systems Handbook*, pp. 1–34 (2015)
13. Rossi, S., Barile, F., Galdi, C., Russo, L.: Recommendation in museums: paths, sequences, and group satisfaction maximization. *Multimedia Tools and Applications* **76**(24), 26031–26055 (2017)
14. Xiao, L., Min, Z., Yongfeng, Z., Zhaoquan, G., Yiqun, L., Shaoping, M.: Fairness-Aware Group Recommendation with Pareto-Efficiency. In: *ACM RecSys '17*. pp. 107–115 (2017)