User-Item Group Formation with GroupFinder

Igo Brilhante[1], Jose Antonio Macedo[1],
Franco Maria Nardini[2], Raffaele Perego[2], Chiara Renso[2]

[1] Federal University of Ceará, Brazil – [2] ISTI-CNR, Italy

Abstract. The GroupFinder framework addresses the new problem of recommending the best group of friends with whom to enjoy a given item, e.g., a travel destination or a movie. Given a user, her social network and a recommended item that is relevant for the user, our novel recommendation task tries to maximize: i) the relevance of the recommended item for every member of the group, and ii) the intra-group social relationships. This extended abstract shortly summarize the work in [4]: we introduce the User-Item Group Formation problem, the possible solutions and the recommendation framework that organizes them. We experiment the proposed solutions using four publicly available Location Based Social Network datasets confirming the effectiveness and the feasibility of the proposed solutions.

1 Introduction

Our work starts from the simple observation that some human activities are better enjoyed with travel companions. This shifts the problem from recommending a single item to a single user (as in the traditional cases) to a new paradigm of recommendation that takes into account items and groups of users. Given a user and a recommended item, we want thus to deal with the novel problem of suggesting the "best" group of friends with whom to enjoy the recommended item. Consider for example a user who has been recommended to visit Paris: we want to be able to suggest the travel companions who can join her in visiting Paris. Such group should ideally have interest in visiting Paris and also be friend each other to facilitate the staying together. Thus we need to balance the strength of the group internal friendship with the group members interest in travelling to Paris. Considering this last scenario, we design a recommendation technique suggesting the "best" group of k friends for a pair $< user, item >$ taking into account both the social relations and the preferences of the user and the group. Since this approach focuses on the formation of the group based on an item and a user, we refer to it as *User-Item Group Formation* problem (UI-GF or simply *group formation*). We present two algorithms as possible solutions to this problem and the global recommendation framework that incorporates them. The full version of this paper has been published in [4], where the details of the algorithms and the evaluation are reported.

2 The User-Item Group Formation problem and the GroupFinder Framework

Let U be a set of users and I a set of items, given a user $u \in U$, her social network $\mathcal{S}$ and an item $i \in I$ suggested to u, UI-GF aims at discovering the group of k friends of u which maximizes a measure modeling the "satisfaction" of the group for the recommended item. This measure of satisfaction considers both the mass appeal of the recommended item for every member of the group and the intra-group social relations. Given the relevance $R(u,i)$ of i for u we need first to extend the measure of relevance to pairs of users. This adaptation is obtained exploiting two well-known relevance aggregation methods: Aggregated Voting ($R_{PAV}(u,v,i) = R(u,i) + R(v,i)$) and Least Misery ($R_{PLM}(u,v,i) = min_{z \in \{u,v\}} R(z,i)$) [6, 1].

Since we aim at weighting differently the interest of an item for a pair of users on the basis of their friendship, we introduce the *pairwise satisfaction* function measuring the relevance of a item i for two users u and v, weighted according the "strength" of their friendship $w(u,v)$.

Definition 1 (Pairwise Satisfaction). *Given an item $i \in \mathcal{I}$ and $u, v \in \mathcal{U}$, the pairwise satisfaction of users u and v w.r.t. the item i is defined as $PS(u,v,i) = w(u,v) \cdot R_P(u,v,i)$.*

On the basis of this pairwise satisfaction we build the following User-Item Ego Network.

Definition 2 (User-Item Ego Network). *Given a user u, an item i, and an integer θ, the User-Item Ego Network of u w.r.t i is defined as an undirected weighted graph $\Gamma^{\theta}_{u,i} = (F, E)$ where $F \subseteq \mathcal{U}$ is the set of friends of u at a distance lower than or equal to θ in the original graph $\mathcal{S}_G$, and E is the set of edges weighted by the pairwise satisfaction $PS(\cdot,\cdot,i)$.*

We model the UI-GF problem of finding the "best" group of k friends of user u for item i as the problem of finding the densest k-subgraph over the user-item ego network. In this formulation the densest k-subgraph problem has the objective of finding the subgraph of exactly k users that maximizes the weighted pairwise satisfaction density. In this way, we go to the problem definition that is to select from F a group of k users characterized by strong friendship relations and high interest w.r.t the proposed item i:

Definition 3 (User-Item Group Formation). *Given a user u, an item i, her user-item ego network $\Gamma^{\theta}_{u,i}$, and an integer k, the User-Item Group Formation problem asks to find the subgraph $G_{u,i} = (F_u, E_u)$ of $\Gamma^{\theta}_{u,i}$, $|F_u| = k$ that maximizes the weighted pairwise satisfaction density:*

$$\max_{\forall G_{u,i} \subseteq \Gamma^{\theta}_{u,i}, |F_u| = k} \rho(G_{u,i}) = \frac{2 \cdot \sum_{\forall t,v \in F_u} PS(t,v,i)}{k \cdot (k-1)}$$

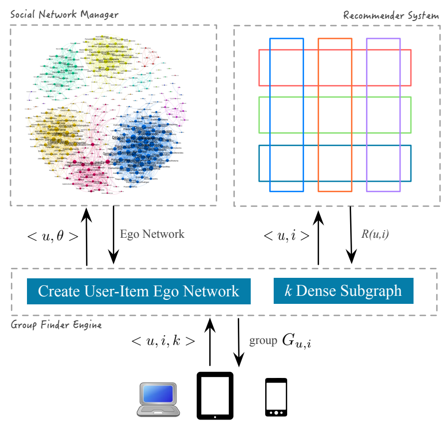

Fig. 1: The GroupFinder framework.

Solving the user-item group formation problem thus requires to compute the densest k-subgraph maximizing the pairwise satisfaction. The densest k-subgraph problem is NP-hard since it generalizes the clique problem [2]. We thus propose a greedy approximation algorithm (Greedy), and a k-Nearest-Neighbor heuristic ($k-$NN). Both these algorithms exploit a measure of pairwise satisfaction aggregated at the level of each user and they are both presented in detail in paper [4].

These algorithms are incorporated into the GroupFinder framework, which includes three different components (see Fig. 1). The **Recommender System** is in charge of providing the relevance $R(u, i)$ of the item i for user u; the **Social Network Manager** retrieves the ego network of focal node u; the **Group Finder Engine** implements the algorithms for approaching the UI-GF [4]. Given a request (u, i, k) it coordinates the interaction aimed at obtaining the user ego network and the relevance scores of item i for all the members of u ego network. Then, it builds the user-item ego network $\Gamma^{\theta}_{u,i}$ by exploiting the pairwise satisfaction function computed for each pairs of users. Finally, the densest k-subgraph is computed and returned as result of the UI-GF instance.

3 Experimental Evaluation

GroupFinder have been compared against state-of-the-art baselines by employing public Location Based Social Networks (LBSN) datasets collected from Foursquare, Brightkite, and Gowalla.

For the experiments we use a content-based recommender system that exploit the metadata associated with venues to measure user-item relevance scores. The relevance score $R(u, i)$ of an item i for a user u is computed as the cosine similarity between the user's preference vector $\boldsymbol{v}_u$ and item's relevance vector

$\boldsymbol{v}_i$ [3]: To evaluate the quality of the groups proposed by GroupFinder we compare them against ground-truth groups, i.e., groups of friends that actually enjoyed a specific venue. We extracted these ground-truth groups from the four datasets looking for sets of users who *checked in at the same place within a fixed temporal window*.

We assess the quality of the group recommended by GroupFinder and the baselines solutions on the basis of different metrics: the weighted pairwise satisfaction density, precision, recall. The first metrics corresponds to Definition 3 and allows us to assess the effectiveness of the algorithms in approximating the densest k-subgraph of the user-item ego network. The used baselines are: (1) **Densest k-Subgraph ($DkSP$)**, a well known algorithm from [5] that aims at selecting the densest k-subgraph from a graph G; (2) **Top k-Nodes** a trivial heuristic to compute the densest k-subgraph without considering the edges.

The results achieved report that the densest k-subgraph-based approaches tend to overcome k-Top. k-Top considers only the user interest measured by $R(\cdot,\cdot)$, thus it may generate groups in which the members are not actually friends in the social network. This explains the lower weighted density obtained with k-Top. Interestingly, $DkSP$ performs remarkably better with PAV than with PLM pairwise user-item relevance. Greedy and k-NN algorithms outperform $DkSP$ and k-Top in terms of weighted density with both PAV and PLM pairwise user-item relevance. Greedy outperforms $DkSP$ from 6% to 17% for PAV, and from 26% to 46% for PLM. This means that it suggests groups characterized by a good balance between friendship and users' relevance, avoiding to include users who are not interested in the item or users that are not well-connected with the rest of the group members. A similar behavior is confirmed when evaluating the performance of the algorithms by using the recall metric since Greedy and k-NN achieve higher recall figures when PLM is used. For the precision metric for the Brightkite and Gowalla datasets, Greedy and k-NN are always the best group formation approaches regardless the pairwise user-item relevance function used. The relatively high values of precision and recall achieved by our solutions (extensively discussed in [4]) demonstrate that they are indeed able to suggest meaningful and relevant groups of friends with whom to enjoy a given venue.

References

1. S. Amer-yahia, B. Omidvar-tehrani, S. B. Roy, and N. Shabib. Group Recommendation with Temporal Affinities. *Proc. EDBT*, pages 421–432, 2015.
2. Y. Asahiro, K. Iwama, H. Tamaki, and T. Tokuyama. Greedily Finding a Dense Subgraph. *J. of Algorithms*, 34(2):203–221, 2000.
3. I. Brilhante, J. A. Macedo, F. M. Nardini, R. Perego, and C. Renso. Where shall we go today?: Planning touristic tours with tripbuilder. In *Proc. CIKM*, 2013.
4. I. Brilhante, J. A. Macedo, F. M. Nardini, R. Perego, and C. Renso. Groupfinder: an item-driven group formation framework. In *Proc. MDM 2016*. IEEE, 2016.
5. U. Feige, G. Kortsarz, and D. Peleg. The dense k-subgraph problem. *Algorithmica*, 29:2001, 1999.
6. A. Jameson and B. Smyth. Recommendation to groups. *The Adaptive Web*, 2007.