Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 98 (2016) 431 – 436

Procedia Computer Science

International Workshop on Data Mining on IoT Systems (DaMIS16)

Towards a Collaborative Filtering Framework for Recommendation in Museums: from Preference Elicitation to Group's Visits

Silvia Rossi[a,*], Francesco Barile[b], Davide Improta[a], Luca Russo[a]

[a]Dipartimento di Ingegneria Elettrica e delle Tecnologie dell'Informazione, University of Naples Federico II, via Claudio, 80125, Naples, Italy
[b]Dipartimento di Matematica e Applicazioni, University of Naples Federico II, via Cintia, 80126, Naples, Italy

Abstract

Recommendation systems based on collaborative filtering methods can be exploited in the context of providing personalized artworks tours within a museum. However, in order to be effectively used, several problems have to be addressed: user preferences are not expressed as rating, items to be suggested are located in a physical space, and users may be in a group. In this work, we present a general framework that, by using the Matrix Factorization (MF) approach and a graph representation of a museum, addresses the problem of generating and then recommending an artworks sequence for a group of visitors within a museum. To reach a high-quality initial personalization, the recommendation system uses a simple, but efficient, elicitation method that is inspired by the MF approach. Moreover, the proposed approach considers the individual or the aggregated artworks' ratings to build up a solution that takes into account the physical location of the artworks.

Peer-review under responsibility of the Program Chairs

Keywords: Cultural Heritage, Sequence Recommendations, Collaborative Filtering, Elicitations

1. Introduction

When it comes to museums, curators design exhibitions with a linear narrative and wearable audio guides are optionally available to provide information to the visitors. This setting is generally very static and the user has to decide whether to follow the suggested paths or to autonomously select the artworks to visit[1]. In this context, Recommendation Systems (RSs) are a natural and technologically ready solution for providing customized tours. Several current attempts are trying to design systems to provide to tourists personalized visiting paths as generated by an RS[2,3,4]. However, employing RS in this scenario is still challenging due to (1) the physicality of the domain, (2) the user profiling and cold start problem, and (3) the presence of groups of people visiting together the museum.

In this work, we presented a general framework of a recommendation system to be used in a museum. Our starting hypothesis is that collaborative filtering approaches, as the Matrix Factorization, can be deployed within museum environments to provide the expected user's ratings on new artworks starting from a preference elicitation

* Corresponding author. Tel.: +39-081-679963.
E-mail address: silvia.rossi@unina.it

Peer-review under responsibility of the Program Chairs
doi:10.1016/j.procs.2016.09.067

process that requires only few interaction steps with the user. Moreover, the Internet of Things (IoT) vision of smart museums, could allow to dynamically update the user preferences from implicit behavior analysis by localizing a visitor with respect to artworks [1]. Artworks viewing times can be mapped into a common user-item ratings matrix, but in order to obtain high-quality initial recommendations, an RS must be provided with an efficient and effective process for gathering information about new users. The proposed approach aims at recommending a sequence of artworks physically located in the space, both in the case of a single visitor and of groups. In particular, the satisfaction that a group of visitors will have as a result of a sequence depends, in part, on how the artworks were ordered during the path. Hence, the proposed approach considers the individual or the aggregated artworks' ratings [5,6] while building up the solution in order to predict the satisfaction of each visitor at the end of the visit. The proposed approach is tested on prototype implementation for a simulated museum environment.

2. A Collaborative Filtering Approach for Museum Tour Recommendation

Typically, in museums, recommendation systems are modeled as content-based filtering systems [7] that suggest similar items to those already appreciated in the past by the user. However, those approaches require an effort in providing the semantic information associated with each artwork, but also the explicit acquisition of information necessary to build the user profile. Instead, when such preferences are obtained online by the user visiting style, there is the risk to recommend only things that are similar to the viewed ones.

On the other hand, approaches that are based on Collaborative Filtering requires a user-item matrix of the ratings expressed by other users on the artworks collection in order to provide recommendations. Our starting hypothesis is that in the IoT vision of smart museums, such matrix can be obtained by localizing visitors with respect to physical objects [1], and so artworks viewing times can be, potentially, mapped into a user-item ratings matrix. With this assumption, collaborative filtering approaches can be deployed within museum environments.

Here, we rely on Matrix Factorization (MF) [8] for providing recommendations. The idea behind MF is the existence of some latent factors that determine how a user will rate an item. This method has become very popular in recent years, as it ensures high scalability and accuracy, also it guarantees a lot of flexibility to adapt to different real situations. In its basic model, the MF models map both the users and the items with an array of f latent factors, so that the user-item interaction is modeled as a scalar product in this space. Therefore, each item i is associated with a vector $\overline{q}_i \in \Re^f$, and each user u is associated with $\overline{p}_u \in \Re^f$. For a given item i (or user u), the j-*th* component of $\overline{q}_i$ (or $\overline{p}_u$) measures how the item (user) has that particular factor j. This quantity can be positive or negative.

2.1. Preference Elicitation

The design of an elicitation method implies making decisions that affect both the effort required to the user for expressing his/her preferences and the recommendations' accuracy. One of the most common used approach in RSs to generate a user preference profile is to rely on the new user to express his/her preferences by rating a fixed number of items. However, this is not an efficient way to convert the workload carried by the user in recommendations as it involves a high user's cognitive effort. It is usually more desirable to start offering the recommendations to the visitors as soon as possible, hence minimizing intrusiveness for the users.

The proposed elicitation method is inspired by the MF algorithm as proposed in [9], since it does not rely on specific metadata that for a museum may not be available. The aim is to use some selected latent factors to create few items sets for which the user have to express his/her preferences (See Figure 1b). The sets are formed by items (called seeds) which represent low and high values of a specific latent factor f. This method consists of the following steps: firstly, to extract f latent factors from the user-item matrix and to assign to each item the vector $\overline{q}_i \in \Re^f$ in the f-dimensional space; then to choose the seeds to display to the user. After a number of interactions between the user and the system, the vector representing the user u preference is placed in f-dimensional space of latent variables.

For each considered latent factor, the seeds selection is carried out on the basis of three criteria:

- *Popularity*: in order to ensure that the user is able to vote the proposed items, only popular items have to be selected (e.g., items that were rated by many users).

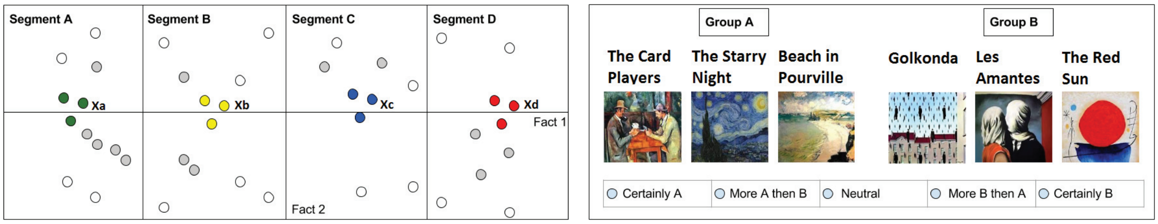

(a) Selection of appropriate items. (b) Interface of the latent factor based elicitation.

Fig. 1: A schematic example of the elicitation process is shown. (a) For each considered factor, items not popular enough (shown in grey) are ignored. The item space is divided into four segments for the currently presented factor. Then, the items near the average values are selected, ensuring that their other characteristics are as neutral as possible with respect to the other factors. (b) Some items from A and D segments are shown to the user to be voted.

- *Isolation*: for each latent factor f, the item space is divided into four segments: A, B, C and D (see Figure 1a). For each segment, a vector is constructed ($\overline{X_A}$, $\overline{X_B}$, $\overline{X_C}$, and $\overline{X_D}$). In such vectors, for the component that represents the considered latent factor, the average items value in the segment (respectively, A, B, C, and D) is calculated. For all other components, such average is calculated on the overall dataset. To ensure that the selected items are good representatives of a particular latent factor and as neutral as possible with respect to other factors, the items near the average values are selected.
- *Diversity*: the items sets displayed to the user are selected so that both sets are easily distinguishable about a particular latent factor (e.g., items from segment A and D). Moreover, too extreme values on one factor may distort the recommendation process, so items that have values above or below the fifth percentile for this particular factor should be removed.

At the end of the seed extraction process f sets of two groups of artworks are shown to the user. Given two items sets A and B, several possibilities are allowed: 1) to express a strong preference for the items set A; 2) to express a strong preference for the items set B; 3) to express more preferences for the items set A then B; 4) to express more preferences for the items set B then A; 5) to express a neutral preference (See Figure 1b). In the latter case, the user does not want or cannot decide because he/she does not have sufficient knowledge of the proposed items. This translates into giving a value to the latent factor in question lower or higher depending on the expressed preference. After each user choice, a vector $\overline{p}_u$, which represents target user interests, is updated. Given that X_A, X_B, X_C, X_D values indicate the average values of the items used to describe low, low-medium, medium-high and high values of factor f, if the user prefers the items set with low value for latent factor f, then f-th component of its vector is set to X_A, else if the user partially prefers the item set with low value then f-th component of its vector is set to X_B, else if the user partially prefers the item set with high value then f-th component of its vector is set to X_C, at the end, if the user prefers the item set with high value, then f-th component of its vector is set to X_D. If the user prefers not to make a decision, the correspond dimension is left blank (also called neutral component).

2.2. Artworks Sequence Generation

Once that the individual user preferences are acquired, the MF approach allows predicting the user rating on the museum collection. In order to take into account the physical disposition of the artworks, the proposed approach relies on the representation of the museum (artworks and paths) as a Directed Acyclic Graph (DAG) that contains possible directed paths from the museum entrance towards the exit.

In a museum scenario, we might have to deal with two possible situations. In the first case, the museum may already have an acyclic default structure of possible paths that are defined by the exhibit designer. Hence, the recommendation system can simply operate on the provided graph. On the contrary, in the second case, we can assume that the graph of the artworks is fully connected, so it is possible to go to every artwork starting from another one. Starting from this representation, it is necessary to make an additional step to transform it in a DAG.

Moreover, the recommendation system should suggest an artworks sequence that maximizes visitors' satisfaction. Therefore, the solution that the system must find can be modeled as a Longest Path problem on a graph that connects the artworks within a museum starting from an entry point towards an exit. However, such Max Satisfaction Path could be bounded by the time available for the museum visit. Here, we focus only on the traveled distance. Hence, the problem of providing recommendations within the museum can be modeled as a Resource Constrained Longest Path Problem (RCLPP). The RCLPP has been proven to be difficult to solve (NP-hard) [10]. However, in the literature, algorithms that solve the problem for DAG in pseudo-polynomial time have been proposed [11,12] and can be used to compute the solution.

In order to model the user satisfaction, we have to take into account that, with respect to the classical view of an RS where the item satisfaction can be directly associated with its rating, here the satisfaction that a group of visitors will have as a result of a sequence of items recommendation depends, in part, on how the artworks were ordered during the path. Following the work of [13], we decided to model the individual user u satisfaction for the artwork i in a sequence as follows:

$$sat_u(P + \langle i \rangle) = \frac{\gamma \times sat_u(P) + r_{u,i}}{1 + \gamma} \tag{1}$$

where $sat_u(P)$ is the user satisfaction of all the previously seen artworks on the path P and $r_{u,i}$ is the estimated rating for the user u with respect to the artwork i as evaluated by the MF. The constant γ is to modulate the contribution of the previous artworks in the evaluation of a specific one i. With $\gamma = 1$ no decrease in satisfaction over time occurs (i.e., the satisfaction of visiting the artwork i depends on all the other seen artworks), and with $\gamma = 0$ no memory of past items would be used (i.e., the satisfaction of visiting the artwork i depends only on its rating).

Finally, since an RS for museum should consider also the presence of groups, an aggregated value that takes into account each member of the group has to be considered. In general, within the group recommender systems literature, the proposed approaches could be divided into two main categories, the *merging preferences*, in which single user preferences are aggregated in order to create a single group profile, on which apply an individual recommendation system, and the *merging recommendations* approach, that consists of aggregating the single recommendations obtained for each user using techniques known as *Social Choice functions* [14]. Since the considered problem involves the computation of a single path for the whole group, we decided to use a merging preferences approach, whereas the satisfaction of seeing the artwork i for group G is then computed using a multiplicative aggregation strategy, e.g., by multiplying the satisfaction of each member of the group to see the artwork i:

$$sat_G(P + \langle i \rangle) = \prod_{u \in G} sat_u(P + \langle i \rangle) \tag{2}$$

3. Offline Evaluations

To evaluate the effectiveness of proposed recommendation general framework, we conducted an offline study on a simulated museum room, composed of 20 locations in which to place the artworks. Since the MF algorithm relies on the availability of the user-item ranking matrix, we decided to use a movie dataset[1] to simulate ratings of artworks. The chosen movies include some popular and some less popular, covering all available genres in the dataset. Then, we set the distances between any two points of interest to generate the expanded DAG. The MF algorithm used to provide recommendation is the *algorithm python-recsys*[2].

3.1. Effectiveness of the preference elicitation

In order to get a good trade-off between the initial effort and the obtained performances, we empirically decide to submit 6 dialog boxes for initial preference profiling corresponding to a 6 dimensions latent factors space. The performance of the elicitation process is evaluated with respect to a baseline, in terms of nDCG and F-measure. As

[1] http://grouplens.org/datasets/movielens/
[2] https://github.com/ocelma/python-recsys.

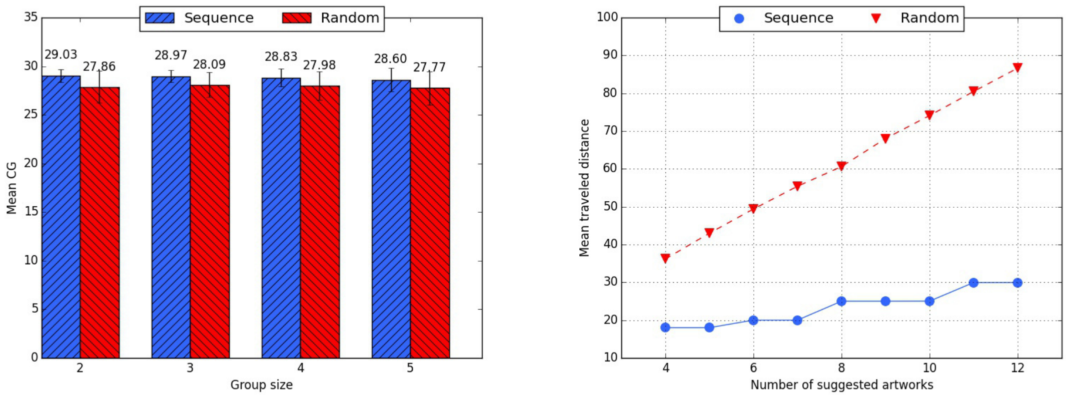

(a) CG mean, W = 20 meters.

(b) Traveled distance mean per number of artwork suggested.

Fig. 2: Sequence vs Random results for groups.

a baseline method, we consider the classical approach where the user has total freedom to express his/her initial preferences by voting at least 20 items.

We conduct a 5-fold cross-validation on the dataset, where, the 80% of the users and their ratings are used to build up the item groups to be shown to the users. The remaining 20% is used to simulate the users decisions with respect to the elicitation process and to evaluate the proposed recommendations. Moreover, for each user in 20% of the test set, 70% of his/her ratings were considered to simulate the initial preference elicitation process (or the initial user profile for the baseline case), while the other 30% was used to evaluate the effectiveness of the recommendations.

For the offline testing, we simulated the behavior of the users starting from the dataset information. In particular, the value (out of five possible choices) to assign to a latent factor was chosen by evaluating the average of the items ratings which represent the two opposite values of the latent factor (e.g., items in the sectors A and D). Subsequently, the difference r of the two values is considered, and it is used to simulate the user's choice (A if $r < -1.5$, B if $-1.5 \leq r < -0.5$, 0 if $-0.5 \leq r < 0.5$, C if $0.5 \leq r < 1.5$ and D if $r \geq 1.5$). For example, consider the scenario where the user u expressed votes for the following items: $m_1 = 3.5, m_2 = 3.5, m_3 = 5$ and $m_4 = 0, m_5 = 3, m_6 = 2$ and, the item $1, 2, 3$ represent low values for the latent factor $fact_1$, while items $4, 5, 6$ represent the high values. Then, since the means of rated items are 4 and 1.66 the simulation assigns to the latent factor a vote in A.

In conclusion, test about nDCG and F-measure shows that there are no substantial differences between proposed approach and baseline method. Indeed, the values of F-measure are for the proposed approach and for the baseline method respectively 37% and 36% and, for nDCG we have 0.183 and 0.182 respectively. In both cases, differences with respect to the baseline case are small. As confirmed by the ANOVA tests, the values are significant with $p < 0.01$.

We found that the exploitation of latent factors is a promising way to generate interactive, choice-based recommendation dialogs. In comparison to baseline method, we achieve a substantive equality with respect to accuracy and user satisfaction about recommendations without sacrificing interaction efficiency.

3.2. Effectiveness of the sequences

To evaluate the proposed sequences, we used the Cumulative Gain (CG). Since the dataset that we used does not contain configurations of real groups, it was necessary to create artificial groups. Two types of groups were created: groups with high inner similarity and random groups. In the first case, the group contains only users who have similar tastes, e.g. groups that provided similar ratings, whereas in the second case the group may also contain users who have very different tastes because the choice was made randomly. In addition, for each type of group, groups of 2, 3, 4 and 5 users have been created.

Our initial hypothesis is that the results obtained with our approach, called *Sequence*, are not due to chance and, for this reason, we compare them with the results obtained with an approach, called *Random*, which suggests artworks at random. A sequence suggested with the *Random* approach contains r artworks, where r depends on the number of artworks suggested by the *Sequence* approach. In our experiment, after the generation of sequence I^* with the *Sequence* approach, random $r = |I^*|$ movies are chosen, without taking into account maximum walking distance threshold W.

In Figure 2a is shown the histogram of the CG mean, for high inner similarity groups of the results obtained with the two approaches, *Sequence* and *Random*. Notice that our approach performs slightly better than the one that suggests artworks randomly. However, the *Random* approach is a bad recommendation strategy of sequences because it does not take into account the traveled distance, forcing visitors to walk too much (Figure 2b).

4. Conclusions

In this work, we presented a general framework for collaborative filtering-based recommendations in museums. It starts with a simple and efficient elicitation method and ends with a recommendation of artworks sequences for groups of visitors. The effectiveness of our framework is evaluated with offline testings in a simulated museum environment. The encouraging results will be more deeply analyzed in the future. First, the aim is to test our system in a real museum with real groups of visitors, and with different satisfaction functions and algorithms for generating sequences. Finally, we would like to consider a temporal threshold of the whole visit that includes also the artworks visiting time and to dynamically optimize the tour according to crowding around artworks.

References

1. W. R. van Hage, N. Stash, Y. Wang, L. Aroyo, The Semantic Web: Research and Applications: 7th Extended Semantic Web Conference, ESWC, Proceedings, Part I, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, Ch. Finding Your Way through the Rijksmuseum with an Adaptive Mobile Museum Guide, pp. 46–59.
2. F. Bohnert, D. F. Schmidt, I. Zukerman, Spatial processes for recommender systems, in: Proceedings of the 21st International Jont Conference on Artifical Intelligence, IJCAI'09, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2009, pp. 2022–2027. URL `http://dl.acm.org/citation.cfm?id=1661445.1661768`
3. I. Roes, N. Stash, Y. Wang, L. Aroyo, A personalized walk through the museum: The chip interactive tour guide, in: CHI '09 Extended Abstracts on Human Factors in Computing Systems, CHI EA '09, ACM, New York, NY, USA, 2009, pp. 3317–3322.
4. F. Amato, A. Chianese, A. Mazzeo, V. Moscato, A. Picariello, F. Piccialli, The talking museum project, Procedia Computer Science 21 (2013) 114 – 121, the 4th International Conference on Emerging Ubiquitous Systems and Pervasive Networks (EUSPN-2013) and the 3rd International Conference on Current and Future Trends of Information and Communication Technologies in Healthcare (ICTH).
5. A. Caso, S. Rossi, Users ranking in online social networks to support pois selection in small groups, in: Posters, Demos, Late-breaking Results and Workshop Proceedings of the 22nd Conference on User Modeling, Adaptation, and Personalization, Vol. 1181 of CEUR Workshop Proceedings, CEUR-WS.org, 2014.
6. S. Rossi, A. Caso, F. Barile, Combining users and items rankings for group decision support, in: Trends in Practical Applications of Agents, Multi-Agent Systems and Sustainability, Vol. 372 of Advances in Intelligent Systems and Computing, Springer International Publishing, 2015, pp. 151–158.
7. P. Lops, M. de Gemmis, G. Semeraro, Recommender Systems Handbook, Springer US, Boston, MA, 2011, Ch. Content-based Recommender Systems: State of the Art and Trends, pp. 73–105.
8. Y. Koren, R. Bell, C. Volinsky, Matrix factorization techniques for recommender systems, Computer 42 (8) (2009) 30–37.
9. B. Loepp, T. Hussein, J. Ziegler, Choice-based preference elicitation for collaborative filtering recommender systems, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM, 2014, pp. 3085–3094.
10. H. Sato, J. O. Royset, Path optimization for the resource-constrained searcher, Naval Research Logistics (NRL) 57 (5) (2010) 422–440.
11. M. Desrochers, F. Soumis, A generalized permanent labeling algorithm for the shortest-path problem with time windows, Infor 26 (3) (1988) 191–212.
12. S. Irnich, G. Desaulniers, et al., Shortest path problems with resource constraints, Column generation 6730 (2005) 33–65.
13. J. Masthoff, A. Gatt, In pursuit of satisfaction and the prevention of embarrassment: affective state in group recommender systems, User Modeling and User-Adapted Interaction 16 (3-4) (2006) 281–319.
14. S. Rossi, F. Cervone, Social utilities and personality traits for group recommendation: A pilot user study, in: Proceedings of the 8th International Conference on Agents and Artificial Intelligence, 2016, pp. 38–46.