

Recommender Systems and Social Networks: an application in Cultural Heritage

Flora Amato, Vincenzo Moscato, Antonio Picariello, Giancarlo Sperli

Dip. di Ingegneria Elettrica e Tecnologie dell'Informazione

University of Naples "Federico II", Naples, Italy

{flora.amato, vmoscato, picus, giancarlo.sperli}@unina.it

Abstract

In the last decade Recommender Systems have become useful tools helping users to find "what they need" from considerable amount of data. One of the more obvious applications of such systems in the Cultural Heritage domain is to assist users when visiting cultural environments (such as museums, archaeological sites, old town centers and so on), providing a multimedia guide that is able to dynamically suggest relevant information available in multiple web repositories (e.g. multimedia sharing systems and on-line social networks). In this paper, we propose a novel recommendation approach that combines several aspects of users - i.e. their preferences (usually in the shape of items' metadata) and interactions within a social community modeled using hypergraphs - together with items' multimedia features and context information within a general framework that can support different applications (touristic guiding services for museums, visiting paths recommendation for old town centers and archeological sites, etc.). Preliminary experiments on user satisfaction show how our approach provides very promising and interesting results.

1 Introduction

The development and promotion of worldwide Cultural Heritage using Information and Communication Technologies (ICT) represent nowadays an important research issue with a variety of potential applications.

In the last decade, such technologies have radically changed the purpose of Cultural Heritage exhibitions that is rapidly moving from an old vision, providing a tourist with static information consisting of a large amount of cultural signs, to novel *personalized services*, matching the visitors' personal goals and behaviors by

considering their cultural needs and preferences and context information.

Indeed, users' experience could be surely enhanced if, instead of using classic "tourist" devices, they could be embedded in a cultural environment with a number of functionalities for representing the relevant information derived from the available digital sources, such as text descriptions, pictures, and videos. In this way, tourists would be given the opportunity of enjoying multimedia stories in real time, thus enriching their cultural knowledge.

From the other hand, we are assisting to an explosive and amazing increase of digital information, and as a consequence, more and more huge data collections of different nature are widely available and have constrained users necessarily to deal with this ocean of information to find "what they need". In particular, on-line social networks (e.g. Facebook) and multimedia sharing systems (e.g. YouTube, Flickr, Panoramio, Instagram, etc.), together with open digital libraries and archives (e.g. DBpedia), constitute the main multimedia information sources that can be considered "useful" for tourists when they visiting cultural environments such as museums, archaeological sites, old town centers and so on.

As well known, *Recommender Systems* have been introduced to facilitate the browsing of such collections, thus realizing the transition in the Web from the *search* to the *discovery* paradigm.

Generally, recommender systems help people in retrieving information that match their preferences by recommending products or services from a large number of candidates, and support people in making decisions in various contexts: what items to buy, which movie to watch, which music to listen, what travels to do, or even who they can invite to their social network, just to make some examples [24, 25].

One of the more obvious applications of such systems in the Cultural Heritage domain is to assist users

when visiting cultural environments, providing a *multimedia guide* that is able to dynamically suggest relevant information available in multiple web repositories.

Formally, a recommender system deals with a set of *users* $U = \{u_1 \dots, u_m\}$ and a set of *items* $O = \{o_1, \dots, o_n\}$. For each pair (u_i, o_j) , a recommender can compute a *score* (or a *rank*) $r_{i,j}$ that measures the expected interest of user u_i in item o_j (or the expected utility of item o_j for user u_i), using a *knowledge base* and a *ranking* algorithm that generally could consider different combinations of the following characteristics: (i) user preferences and past behavior, (ii) preferences and behavior of the user community, (iii) items' features and how they can match user preferences, (iv) user feedbacks, (v) context information (i.e. user location, observed items, weather and environmental conditions, etc.) and how recommendations can change together with the context.

In the literature, surveys on recommender systems usually classify the different kinds of approaches in four main categories: *content-based* [22, 29, 30] (with their extensions to deal multimedia data and their features [19, 13, 20]), *collaborative filtering* [2, 29, 23, 17] (with their customizations to take into account social elements as user reviews and opinions [32, 28, 18, 27, 11, 21, 9]), *hybrid* [26] and *context aware* [10, 15, 16] techniques. Finally, a recent category of recommenders, named *Large Scale Recommender Systems* (LSRS) [31], calls for new capabilities of such applications to deal with very large amount of data with respect to scalability and efficiency issues.

In our opinion, modern recommending applications have to take into account in some way all the above characteristics to provide useful and reliable recommendations both for virtual and physical environments. To this goal, the last generation of recommender systems is usually composed by one or more of the following components [25].

A *pre-filtering* module that selects for each user u_i a subset $O_i^c \subset O$ containing items that are good candidates to be recommended; such items usually match user preferences and needs.

A *ranking* module that assigns w.r.t. user u_i a rank $r_{i,j}$ to each candidate item o_j in O_i^c using the well-known recommendation techniques (i.e., *content-based*, *collaborative filtering* and *hybrid* approaches) that can exploit in several ways items' features and users' preferences, feedbacks (in the majority of cases in terms of *ratings*) and behavior.

A *post-filtering* module that dynamically excludes, for each user u_i , some items from the recommendations' list; in this way, a new set $O_i^f \subseteq O_i^c$ is obtained on the base of user feedbacks, other contextual information

(such as data coming from the interactions between the user and the application) and possible additional constraints.

In this paper, we propose a novel recommendation approach that combines several aspects of users - i.e. their *preferences* (in the shape of items' metadata) and *interactions* (user to user and user to content) within a social community modeled using hypergraphs - together with items' *multimedia features* and *context information* within a general framework that can support different applications (touristic guiding services for museums, visiting paths recommendation for old town centers and archeological sites, etc.).

In other words, it is the user with his/her preferences (in the pre-filtering stage) and actions (in the post-filtering stage) to drive the recommendation process towards the real useful items among those that a social community considers the "best ones" (computed in the ranking stage), as in a collaborative filtering approach, where a user "learns by the others" the item utility, on the base of an *influence* measure.

The paper is organized as follows. Section 2 provides a functional overview of our system and describes the proposed strategy for recommendation. Section 3 illustrates a system customization for a tourist multimedia guide, reporting some implementation details. Section 4 reports preliminary experimental results, and provides a comparison with other recommendation techniques. Finally, Section 5 gives some concluding remarks and discusses future work.

2 The framework

2.1 System Overview

Figure 1 describes at a glance an overview of the proposed system.

Multimedia data to be recommended are retrieved by a *Wrapper* component that is composed by several modules. The *Crawler* is responsible of: (i) periodically accessing to the items' repositories (e.g., *Instagram*, *Flickr*, *Panoramio*, *Google Images*, *YouTube*, *Facebook*, *DBpedia* etc.), (ii) extracting for each item all the *features* (e.g., metadata, multimedia descriptions, etc.) and other information (e.g. user preferences, comments, time-stamped items' observations and all the different interactions between users and objects). A part of such information will be then exploited by the *Hypergraph Learning* module to build the *hypergraph* modeling the entire *Multimedia Social Network* (MSN)[6]. After the wrapping phase, all the information are stored in the *Knowledge Base* of the

system. In particular, it is composed by: (i) the *Multimedia Social Network Hypergraph*, (ii) the *Items DB* containing items with all the related features, (iii) *User Profiles* containing user preferences, (iv) *Contextual Data* containing some additional context information (e.g. user location, weather conditions, etc.).

Multimedia items are then grouped by the related *cultural Points Of Interest* (POIs): e.g. paintings of museum rooms, buildings in ancient ruins or in an old town center, etc.

The *Recommender Engine* provides a set of recommendation facilities for multi-dimensional and interactive browsing of items. Exploiting user preferences, the *Prefiltering* module selects a set of *candidate* items for recommendation; successively, the *Objects Ranking* module assigns a ranking of such candidates exploiting some *ranking functions* defined on the MSN. Finally, the *Postfiltering* module dynamically selects on the base of some constraints (e.g. the item that a user is currently watching and context information) a subset of candidates.

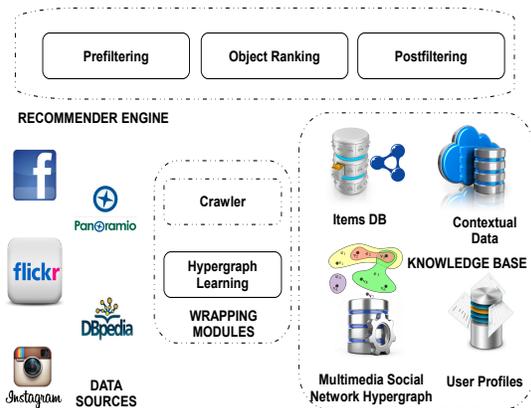


Figure 1. System Overview.

2.2 Recommendation Process

2.2.1 Pre-filtering Stage using user preferences

In the *pre-filtering* stage, our aim is to select for a given user u_h a subset $O_h^c \subset O$ containing items that are good “candidates” to be recommended.

Each item subjected to recommendation may be represented in different and heterogeneous feature spaces. For instance, a picture may be described by a set of

metadata as title, description, tags, by the position in which was taken and so on. Each of these sets of features contributes to the characterization of the items to different *extents*.

The first step consists in clustering together “similar” items, where the similarity should consider all (or subsets of) the different spaces of features. To this purpose, we employ *high-order star-structured co-clustering* techniques - that some of the authors have adopted in previous work [14, 7, 8] - to address the problem of heterogeneous data pre-filtering.

Let $O = \{o_1, \dots, o_n\}$ be the set of items and $\mathcal{F} = \{F^1, \dots, F^l\}$ a set of l feature spaces. In our recommendation problem, a user u_h is represented as a set of vectors in the same l feature spaces describing the items. To provide a first candidate list of items to be recommended, we measure the *cosine distance* of the user vectors associated to the k -th space, with the centroids of each item clusters in the k -th space. For each space, the most similar item cluster is chosen leading to l clusters $\{X_1^c, \dots, X_l^c\}$ of candidate items.

Then, two different strategies can be adopted to provide the pre-filtered list of candidate items O_h^c : (i) *set-union strategy* - the items belonging to the union of all clusters are retained, i.e., $O_h^c = \bigcup_k X_k^c$; (ii) *threshold strategy* - the items that appears in at least ths clusters ($ths \in \{1 \dots l\}$) are retained.

2.2.2 Ranking Stage via hypergraph modeling

The main goal of this stage is to automatically rank the set of items O embedding in a collaborative learning context: the MSN deriving by the integration of the different multimedia data sources. In particular, we use a novel technique that the authors have proposed in a previous work [6]. In our vision, a MSN is basically composed by three different kinds of entities (nodes):

- *Users* - the set of persons and organizations constituting the particular social community: several information concerning their profile, interests, preferences, etc. can eventually be considered and exploited by our model;
- *Multimedia Objects* - the set of multimedia resources (i.e. images, video, audio, posts, documents, etc.) that can be shared within a MSN community: high level (*metadata*) and low level information (*features*) can be properly used in our model;
- *Annotation Assets* - each set of symbols (e.g., keyword, tag, label, etc.) exploited by users to annotate multimedia resources within a MSN; we explicitly note that it is possible to relate a given as-

set with a specific *concept* (as an example a topic, a named entity, etc. which definition can be found into dictionaries, ontologies and so on), thus formally providing the related semantics.

Several types of relationships can be established among the described entities: a user can annotate an object with a particular tag, two friends can comment the same object, a user can tag another user in a photo, a user can share an object within a group, etc. In particular, we distinguish between *user to user relationships*, describing user actions towards other users, and *user to multimedia relationships*, describing user actions on objects, eventually involving some annotation assets. In addition, *similarity relationships* can be added between two objects (using multimedia features) or between two assets (by taxonomic distances).

Due to the variety and complexity of these relationships, we leverage the *hypergraph* formalism to model a MSN (all the details are provided in [6]). Then, we introduce some functions can be profitably used to “rank” users or multimedia objects in a MSN.

In our model the concept of *rank* of a given node is related to the concept of *influence*, and in our vision it can be measured by the number of user nodes that are “reachable” within a certain number of steps using any hyperpath, with respect to a social community of users, and eventually to a given *topic* of interest.

The final goal is to compute the ranking of the multimedia items in O_h^c , using as measure the social influence of each object within the users’ community.

2.2.3 Post-Filtering Stage by context information

In this stage, we have introduced a *post-filtering* method for generating the final set of “real” candidates for recommendation using *context* information.

The context is represented by means of the well-known *key-value* model [1] using as dimensions some of the different feature spaces related to items. In our system, context features can be expressed either directly using some *target items* (e.g. objects that have positively captured user attention) or specifying the related values in the shape of *constraints* that recommended items have to satisfy.

Assume that a user u_h is currently interested in a target item o_j . We can define the set of candidate recommendations as follows:

$$O_{h,j}^f = \bigcup_{k=1}^M \{o_i \in O_h^c \mid a_{ij}^k > 0\} \cup \{o_i \in NNQ(o_j, O_h^c)\} \quad (1)$$

The set of candidates includes the items that have been accessed by at least one user within k steps from o_j , with k between 1 and M , and the items that are most similar to o_j according to the results of a *Nearest Neighbor Query* ($NNQ(o_j, O_h^c)$) functionality. Note that a positive element a_{ij}^k of A^k indicates that o_i was accessed exactly k steps after o_j at least once. The ranked list of recommendations is then generated by ranking the items in $O_{h,j}^f$, for each item o_j selected as interesting by user u_h , using the ranking vector R_h thus obtaining the final set O_h^f .

Finally, for each user all the items that do not respect possible context constraints are removed from the final list.

3 A Case Study

We have opportunely customized our system in order to provide touristic multimedia guiding services for users that are interested in visiting the old town center of Naples, Italy. On the base of user preferences and actual position, a set of POIs are shown on a proper map to tourists correlated with a multimedia description.

For instance, when a user is approaching a particular cultural POI (e.g. *Piazza del Gesu’ Nuovo*), the related multimedia description and the set of candidate objects (i.e. images and texts related to the near POIs) are delivered on the user’s mobile device (pre-filtering stage).

The list of proposed objects depends on the user’s preferences (e.g. the majority of items will be images if a user prefers to see such kinds of data and will reveal effective user needs), is initially ordered according to effective user location (i.e. the closest items will appear at the top of list) and contains data grouped by the related cultural POI. Successively, after the user has selected one or more objects (for example the item he is currently watching), the recommendation services first perform a ranking (ranking stage) of all the candidate objects according to their recommendation grades and then filters the recommendation list considering only the most similar items to target objects (post-filtering stage).

When a user is near to a different POI, he/she can decide to modify the list of target objects (e.g. removing those related to the previous visited POI or adding new objects) and consequently recommendations will be automatically updated, thus including new items.

The design choices are briefly reported in the following.

- We consider as data source *Flickr*, *Instagram*, *Panoramio*, *DBPedia* and other domain digital

libraries, collecting about 500,000 items (images and texts related to historical buildings, churches, famous square and other attractions) and about 5,000 user profiles.

- As items' metadata, we consider for each multimedia item information related to *title, description, type, kind, language, tags, keyword, comments, ratings* (and for pictures the geographic position in which were taken). In addition, images are also described by a set of low-level features (i.e. SURF).
- For each item, available users' preferences, comments, feedbacks and other actions have been captured, also exploiting correlated public information from Social Networks (i.e. *Facebook*).

For what implementation details concern, the Wrapping modules leverage proper API and JAVA libraries to collect the different information of interest.

The Knowledge Base, realized using different technologies, allows to manage all the different kind of information: Contextual Data instances (messages containing information about users' position) are managed by the *Cassandra* DBMS, Items' descriptions are stored in the *Turtle* format and managed by the *AllegraGraph* repository (semantics of data can be specified by linking values of some attributes to some available ontological schema), User Profiles and the MSN hypergraph are respectively managed by *MongoDB* and *Neo4j* DBMSs.

On the other hand, the Recommender Engine exploits proper JAVA libraries (some developed for the system presented in [5] and integrated with co-clustering libraries [7] and the rank refining procedure¹) to accomplish its tasks.

Finally, a user can interact with our system using at the moment an Android Multimedia Guide App exploiting *Google Map* API.

4 Experimental Results

Recommender Systems are very complex applications that are based on a combination of several models, algorithms and heuristics. This complexity makes evaluation efforts very difficult and thus results are hardly generalizable, as reported in the literature [3]. Moreover, characterizing and evaluating the quality of a user's experience and subjective attitude toward the acceptance of recommender technology is an important issue which we will consider in the following.

¹we use *LIRE* for the content-based image retrieval

The majority of research efforts on recommender system evaluation have mainly focused on prediction *accuracy* and *stability* (e.g., [3]).

More recently, researchers began examining issues related to users subjective opinions and developing additional criteria to evaluate recommender systems. In particular, they suggest that user satisfaction does not always (or, at least, not only) correlate with the overall recommenders accuracy.

Starting from these considerations and based on current trends in the literature, we decided to perform a *user-centric* evaluation based on *user satisfaction* with respect to assigned activities, evaluating how our recommendations can effectively support browsing tasks of different complexity when the complexity of desired items increases.

As in our previous work [4, 5, 7, 8], we evaluate the impact of the proposed system on users engaged in several *search tasks* of multimedia items and compared its performances with the well-known *Panoramio* system² that, in turn, provides basic search mechanisms.

In particular, our goal was to establish how helpful our system is in assisting the search of specific multimedia objects (images) and guiding the users towards information which satisfy their interests. The dataset used in these experiments is a subset of about 10,000 items related to specific POIs.

In order to evaluate the impact of the system on the users, we have conducted the following experiments. The system was made available to a set of 50 users. These users were all interested in the cultural heritage domain, they already had experience in the use of PCs and electronic devices, even if they were not experts in ICT. We asked these users to browse the collection of items and complete several search tasks (20 tasks per user) of different complexity (five tasks for each complexity level), using *Panoramio* facilities. After this test, we asked them to browse the same collection with the assistance of our recommender system and complete other 20 tasks of similar complexity. We have subdivided browsing tasks in the following four broad categories:

1. **Low Complexity** search tasks (T_1): e.g. find at least 30 images related to 3 different POIs depicting ancient churches;
2. **Medium Complexity** search tasks (T_2): e.g. find at least 50 images related to 5 different POIs depicting ancient churches, historical building and famous squares (10 objects for each subject);

²<http://www.panoramio.com/>

Table 1. Comparison between our system and Panoramio in terms of t_a and n_c average values

Task Class	System	$t_a(sec)$	n_c
Low Complexity	Recommender	155	38
Low Complexity	Panoramio	164	42
Medium Compl.	Recommender	335	84
Medium Compl.	Panoramio	402	104
High Complexity	Recommender	1156	298
High Complexity	Panoramio	1302	334
Very High Compl.	Recommender	1645	351
Very High Compl.	Panoramio	1832	410

3. **High Complexity** search tasks (T_2): e.g. find at least 100 images related to 10 different POIs (near to the actual user position) depicting ancient churches, historical building and famous squares (10 objects for each subject);
4. **Very High Complexity** search tasks (T_2): e.g. find at least 150 images objects related to 10 different POIs (near to the actual user position) depicting ancient churches, historical building and famous squares (15 objects for each subject).

Note that the complexity of a task depends on several factors: the number of items to explore, the type of desired features and the number of additional constraints. Two strategies were used to evaluate the results of this experiment: (i) empirical measurements of access complexity in terms of *mouse clicks* and *time*; (ii) TLX (*NASA Task Load Index factor*).

With respect to the first strategy, we measured the following parameters: (i) *access time* (t_a) – the average time spent by the users to request and access all the images for a given class of tasks; (ii) *number of clicks* (n_c) – the average number of clicks necessary to collect all the requested images for a given class of tasks.

Table 1 reports the average values of t_a and n_c for both Panoramio and our system (Recommender), for each of the four task complexity levels defined. Especially for the most complex tasks, our system shows better performances than Panoramio, especially for the more complex tasks.

We then asked the same group of users to express their opinion about the capability of Panoramio and our system respectively to provide an effective user experience in completing the assigned search tasks, based on the TLX evaluation protocol [12].

Specifically, TLX is a multi-dimensional rating procedure that provides an overall workload score based on

Table 2. Comparison between our system and Panoramio in terms of TLX factors for each category of users

TLX factor	Recommender	Panoramio
Mental demand	39	41
Physical demand	36.3	48
Temporal demand	39	50
Effort	35	50.5
Performances	69.7	79.8
Frustration	33.2	44.1

a weighted average of ratings on six sub-scales: mental demand, physical demand, temporal demand, own performance, effort and frustration. Lower TLX scores are better and the average scores are then reported in Table 2.

Our system outperforms in a significative way Panoramio in every sub-scale except for *mental demand* and *performance*: this happens because sometimes an expert user considers the automatic suggestions not useful, just because they know what they are looking for.

In summary, our system provides a better (less frustrating) user experience during the search tasks. In addition, the fact that search tasks can be completed faster using our system is an indication that recommendations are effective, as they allow a user to explore interesting and related items one after another, without the interference of undesired items that would otherwise slow down the process.

5 Conclusions and Future Work

In this paper a novel multimedia and social recommendation approach for Cultural Heritage applications. It combines several aspects of users - i.e. their *preferences* (usually in the shape of items' metadata) and *interactions* within a social community modeled using hypergraphs - together with items' *multimedia features* and *context information* within a general framework that can support different applications (touristic guiding services for museums, visiting paths recommendation for old town centers and archeological sites, etc.).

Preliminary experiments on user satisfaction demonstrated how our approach achieve very promising and interesting results. Future works will be devoted to extend the experimental evaluation to a larger multimedia data set, also considering the performance, evaluated in terms of accuracy, precision

and recall, of the performed recommendations. Moreover, we plan to apply our approach to other kinds of data gathered from heterogeneous collections and compare our approach with other ones proposed in the literature.

References

- [1] Gediminas Adomavicius, Ramesh Sankaranarayanan, Shahana Sen, and Alexander Tuzhilin. Incorporating contextual information in recommender systems using a multidimensional approach. *ACM Transactions on Information Systems (TOIS)*, 23(1):103–145, 2005.
- [2] Gediminas Adomavicius and Alexander Tuzhilin. 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*, 17:734–749, 2005.
- [3] Gediminas Adomavicius and Jingjing Zhang. Stability of recommendation algorithms. *ACM Transactions on Information Systems (TOIS)*, 30(4):23, 2012.
- [4] Massimiliano Albanese, Antonio d’Acerno, Vincenzo Moscato, Fabio Persia, and Antonio Picariello. A multimedia semantic recommender system for cultural heritage applications. In *Semantic Computing (ICSC), 2011 Fifth IEEE International Conference on*, pages 403–410. IEEE, 2011.
- [5] Massimiliano Albanese, Antonio d’Acerno, Vincenzo Moscato, Fabio Persia, and Antonio Picariello. A multimedia recommender system. *ACM Transactions on Internet Technology (TOIT)*, 13(1):3, 2013.
- [6] Flora Amato, Vincenzo Moscato, Antonio Picariello, and Giancarlo Sperli’. Multimedia social network modeling: a proposal. In *Semantic Computing (ICSC), 2016 10th IEEE International Conference on*. IEEE, 2016.
- [7] Ilaria Bartolini, Vincenzo Moscato, Ruggero G Pensa, Antonio Penta, Antonio Picariello, Carlo Sansone, and Maria Luisa Sapino. Recommending multimedia visiting paths in cultural heritage applications. *Multimedia Tools and Applications*, pages 1–30, 2014.
- [8] Francesco Colace, Massimo De Santo, Luca Greco, Vincenzo Moscato, and Antonio Picariello. A collaborative user-centered framework for recommending items in online social networks. *Computers in Human Behavior*, 51:694–704, 2015.
- [9] Ruihai Dong, Michael P. O’Mahony, Markus Schaal, Kevin McCarthy, and Barry Smyth. Sentimental product recommendation. In *Proceedings of the 7th ACM Conference on Recommender Systems, RecSys ’13*, pages 411–414, New York, NY, USA, 2013. ACM.
- [10] Paul Dourish. What we talk about when we talk about context. *Personal and ubiquitous computing*, 8(1):19–30, 2004.
- [11] Gayatree Ganu, Yogesh Kakodkar, and Amélie Marian. Improving the quality of predictions using textual information in online user reviews. *Inf. Syst.*, 38(1):1–15, March 2013.
- [12] Sandra G Hart and Lowell E Staveland. Development of nasa-tlx (task load index): Results of empirical and theoretical research. *Advances in psychology*, 52:139–183, 1988.
- [13] Yoshinori Hijikata, Kazuhiro Iwahama, and Shogo Nishida. Content-based music filtering system with editable user profile. In *Proceedings of the 2006 ACM symposium on Applied computing, SAC ’06*, pages 1050–1057, New York, NY, USA, 2006. ACM.
- [14] Dino Ienco, Céline Robardet, Ruggero G Pensa, and Rosa Meo. Parameter-less co-clustering for star-structured heterogeneous data. *Data Mining and Knowledge Discovery*, 26(2):217–254, 2013.
- [15] Katerina Kabassi. Personalisation systems for cultural tourism. In *Multimedia services in intelligent environments*, pages 101–111. Springer, 2013.
- [16] Alexandros Karatzoglou, Xavier Amatriain, Linas Baltrunas, and Nuria Oliver. Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 79–86. ACM, 2010.
- [17] Yehuda Koren. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining, KDD ’08*, pages 426–434, New York, NY, USA, 2008. ACM.

- [18] Cane WK Leung, Stephen CF Chan, and Fu-lai Chung. Integrating collaborative filtering and sentiment analysis: A rating inference approach. In *Proceedings of The ECAI 2006 Workshop on Recommender Systems*, pages 62–66. Citeseer, 2006.
- [19] Veronica Maidel, Peretz Shoval, Bracha Shapira, and Meirav Taieb-Maimon. Evaluation of an ontology-content based filtering method for a personalized newspaper. In *Proceedings of the 2008 ACM conference on Recommender systems*, RecSys '08, pages 91–98, New York, NY, USA, 2008. ACM.
- [20] Katarzyna Musial, Krzysztof Juszczyszyn, and Przemyslaw Kazienko. Ontology-based recommendation in multimedia sharing systems. *System Science*, 34:97–106, 2008.
- [21] Nikolaos Pappas and Andrei Popescu-Belis. Sentiment analysis of user comments for one-class collaborative filtering over ted talks. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '13, pages 773–776, New York, NY, USA, 2013. ACM.
- [22] Michael Pazzani and Daniel Billsus. Content-Based Recommendation Systems. pages 325–341. 2007.
- [23] Naren Ramakrishnan, Benjamin J. Keller, Batul J. Mirza, Ananth Y. Grama, and George Karypis. Privacy risks in recommender systems. *IEEE Internet Computing*, 5:54–62, November 2001.
- [24] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. GroupLens: An open architecture for collaborative filtering of netnews. pages 175–186. ACM Press, 1994.
- [25] Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor, editors. *Recommender Systems Handbook*. Springer, 2011.
- [26] Andrew I. Schein, Alexandrin Popescul, Lyle H. Ungar, and David M. Pennock. Methods and metrics for cold-start recommendations. In *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '02, pages 253–260, New York, NY, USA, 2002. ACM.
- [27] Vivek Kumar Singh, Mousumi Mukherjee, and Ghanshyam Kumar Mehta. Combining collaborative filtering and sentiment classification for improved movie recommendations. In *International Workshop on Multi-disciplinary Trends in Artificial Intelligence*, pages 38–50. Springer, 2011.
- [28] Johann Stan, Fabrice Muhlenbach, Christine Largeron, et al. Recommender systems using social network analysis: Challenges and future trends. *Encyclopedia of Social Network Analysis and Mining*, pages 1–22, 2014.
- [29] Xiaoyuan Su and Taghi Khoshgoftaar. A survey of collaborative filtering techniques. *Advances in Artificial Intelligence*, 2009, 2009.
- [30] Hilmi Yildirim and Mukkai S. Krishnamoorthy. A random walk method for alleviating the sparsity problem in collaborative filtering. In *Proceedings of the 2008 ACM conference on Recommender systems*, RecSys '08, pages 131–138, New York, NY, USA, 2008. ACM.
- [31] Hsiang-Fu Yu, Cho-Jui Hsieh, Si Si, and Inderjit S Dhillon. Parallel matrix factorization for recommender systems. *Knowledge and Information Systems*, pages 1–27, 2013.
- [32] Xujuan Zhou, Yue Xu, Yuefeng Li, Audun Josang, and Clive Cox. The state-of-the-art in personalized recommender systems for social networking. *Artif. Intell. Rev.*, 37(2):119–132, February 2012.