

Content-Based Personalization Services Integrating Folksonomies

Cataldo Musto, Fedelucio Narducci, Pasquale Lops,
Marco de Gemmis, and Giovanni Semeraro

Department of Computer Science,
University of Bari “Aldo Moro”, Italy
{musto,narducci,lops,degemmis,semeraro}@di.uniba.it
<http://www.di.uniba.it/>

Abstract. Basic content-based personalization consists in matching up the attributes of a user profile, in which preferences and interests are stored, with the attributes of a content object. The Web 2.0 (r)evolution has changed the game for personalization, from ‘elitary’ Web 1.0, written by few and read by many, to web content generated by everyone (*user-generated content* - UGC), since the role of people has evolved from passive consumers of information to that of active contributors.

One of the forms of UGC that has drawn most attention of the research community is folksonomy, a taxonomy generated by users who collaboratively annotate and categorize resources of interests with freely chosen keywords called tags.

FIRSt (**F**olksonomy-based **I**tem **R**ecommender **s**ystem) is a content-based recommender system developed at the University of Bari which integrates UGC (through social tagging) in a classic content-based model, letting users express their preferences for items by entering a numerical rating as well as to annotate rated items with free tags. FIRSt is capable of providing recommendations for items in several domains (e.g., movies, music, books), provided that descriptions of items are available as text documents (e.g. plot summaries, reviews, short abstracts). This paper describes the system general architecture and user modeling approach, showing how this recommendation model has been applied to recommend the artworks located at the Vatican Picture Gallery (Pinacoteca Vaticana), providing users with a personalized museum tour tailored on their tastes.

Keywords: Content-based Recommender Systems, Web 2.0, Folksonomy, Machine Learning, Semantics.

1 Introduction

The amount of information available on the Web and in Digital Libraries is increasing over time. In this context, the role of user modeling and personalized information access is becoming crucial: users need a personalized support in sifting through large amount of retrieved information according to their interests.

Information filtering systems, relying on this idea, adapt their behavior to individual users by learning their preferences during the interaction in order to construct a profile of the user that can be later exploited in selecting relevant items. Indeed, content personalization basically consists in matching up the attributes of a user profile, in which preferences and interests are stored, against the attributes of a content object.

Recently, Web 2.0 (r)evolution has changed the game for personalization, from ‘elitary’ Web 1.0, written by few and read by many, to web content generated by everyone (*user-generated content* - UGC), since the role of people has evolved from passive consumers of information to that of active contributors.

UGC refers to various kinds of media content, publicly available, that are produced by end-users. For example, on Amazon.com the majority of content is prepared by administrators, but numerous user reviews of the products being sold are submitted by regular visitors to the site. Similarly, collaborative platforms like YouTube, Wikipedia, Flickr, Del.icio.us, Photobucket, although supplying heterogeneous content, are completely founded on data freely created by users.

In [3] we began an analysis of UGC impact (namely, tags impact) in recommender systems. The main outcome of that work was that UGC integration in a content-based recommender system causes an increase of the predictive accuracy in the process of filtering relevant items for users.

In this work we will continue this analysis proving how this higher accuracy showed by recommender systems can be exploited for personalization goals in a real-world application. FIRSt (First Item Recommender SysTEM) is a content-based recommender system developed at the University of Bari which integrates UGC (through social tagging) in a classic content-based model, letting users to express their preferences for items by entering a numerical rating as well as to annotate rated items with free tags. FIRSt is capable of providing recommendations for items in several domains (e.g., movies, music, books), provided that descriptions of items are available as text documents (e.g. plot summaries, reviews, short abstracts). We exposed FIRSt basic features through a set of web services, and we exploited them in a museum personalization scenario, letting Vatican Picture Gallery visitors receive suggestions about artworks they could be interested in and tailoring museum tours on their tastes.

This research has been conducted within the CHAT project (Cultural Heritage fruition and e-learning applications of new Advanced multimodal Technologies), that aims at developing new systems and services for multimodal fruition of cultural heritage content. Data has been gathered from the collections of the Vatican Picture Gallery, for which both images and detailed textual information of paintings are available, and letting users involved in the study both rate and annotate them with tags.

The paper is organized as follows. Section 2 introduces the general problem of information filtering and recommender systems; the architecture of FIRSt is described in Section 3 whereas Section 4 focuses the attention on design and development of web services to expose FIRSt functionalities. The experimental

session carried out to evaluate the effectiveness of implemented web services is presented in Section 5. Related work are briefly analyzed in Section 6, while conclusions and directions for future work are drawn in the last section.

2 Information Filtering and Recommender Systems

As proved by the continuous growth of web sites which embody recommender systems as a way of personalizing their content for users, nowadays these systems represent the main field of application of principles and techniques coming from Information Filtering (IF) [9].

As IF techniques usually perform a progressive removal of non-relevant content according to the information stored in a user profile, recommendation algorithms process information about customer interests - acquired in an explicit (e.g., letting users express their opinion about items) or implicit (e.g., studying some behavioural features) way - and exploit this data to generate a list of recommended items, guiding users in a personalized way to interesting or useful objects in a large space of possible options [4].

Among different recommendation techniques that have already been put forward in studies on this matter, the collaborative and content-based filtering approaches are the most widely adopted to date.

Content-based systems analyze a set of documents, usually textual descriptions of the items previously rated by an individual user, and build a model or profile of user interests based on the features of the objects rated by that user [8]. In this approach static content associated to items (the plot of a film, the description of an artwork, etc.) is usually exploited. The profile is then used to recommend new relevant items.

Collaborative recommender systems differ from content-based ones in that user opinions are used, instead of content. User ratings about objects are gathered and stored in a centralized or distributed database. To provide recommendations to user X, the system firstly computes the neighborhood of that user (i.e. the subset of users that have a taste similar to X). Similarity in taste is measured by computing the closeness of ratings for objects that were rated by both users. The system then recommends objects that users in X's neighborhood indicated to like, provided that they have not yet been rated by X.

Although each type of filtering method has its own weaknesses and strengths [1,12], in this work we focused our attention only on a single class of recommenders, introducing in the next section the general architecture of FIRSt, which represents the core of personalization mechanisms designed for Vatican Picture Gallery scenario.

3 FIRSt: Folksonomy-Based Item Recommender system

FIRSt is a semantic content-based recommender system capable of providing recommendations for items in several domains (e.g., movies, music, books), provided

27) Caravaggio - Deposition from the Cross



Painting Description

The Deposition, considered one of Caravaggio's greatest masterpieces, was commissioned by Girolamo Vittrice for his family chapel in S. Maria in Vallicella (Chiesa Nuova) in Rome. In 1797 it was included in the group of works transferred to Paris in execution of the Treaty of Tolentino. After its return in 1817 it became part of Pius VII's Pinacoteca. Caravaggio did not really portray the Burial or the Deposition in the traditional way, inasmuch as Christ is not shown at the moment when he is laid in the tomb, but rather when, in the presence of the holy women, he is laid by Nicodemus and John on the Anointing Stone, that is the stone with which the sepulchre will be closed. Around the body of Christ are the Virgin, Mary Magdalene, John, Nicodemus and Mary of Cleophas, who raises her arms and eyes to heaven in a gesture of high dramatic tension. Caravaggio, who arrived in Rome towards 1592-93, was the protagonist of a real artistic revolution as regards the way of treating subjects and the use of colour and light, and was certainly the most important personage of the "realist" trend of seventeenth century painting.

Popular Tags: caravaggio (5) deposition (5) cross (4) christ (2) vangel (1) maddale (1) unction (1) sepulchre (1) nicodemo (1) virgin (1)

Rate this painting and enter comma separated tags

1 2 3 4 5

Rate this Painting

Fig. 1. Screenshot of Learning Platform

that descriptions of items are available as text documents (e.g. plot summaries, reviews, short abstracts) [6].

In the context of cultural heritage personalization, for example, an artwork can be represented by at least three textual components (called slots), namely *artist*, *title*, and *description*.

The inceptive idea behind FIRSt is to include folksonomies in a classic content-based recommendation model, integrating *static* content describing items with *dynamic* user-generated content (namely tags, through social tagging of items to be recommended) in the process of learning user profiles.

Tags are collected during the training step (Figure 1) by letting users:

1. express their preferences for items through a numerical rating
2. annotate rated items with free tags.

Tags are then stored into an additional slot, different from those containing static content, and are exploited in the profile learning phase in order to include them in the user profiles.

The general architecture of FIRSt is depicted in Figure 2. The recommendation process is performed in three steps, each of which is handled by a separate component:

- **CONTENT ANALYZER** – it allows introducing semantics in the recommendation process by analyzing documents and tags in order to identify relevant concepts representing the content. This process selects, among all the possible meanings (senses) of each polysemous word, the correct one according to the context in which the word occurs. In this way, documents and tags are represented using concepts instead of keywords, in an attempt to overcome

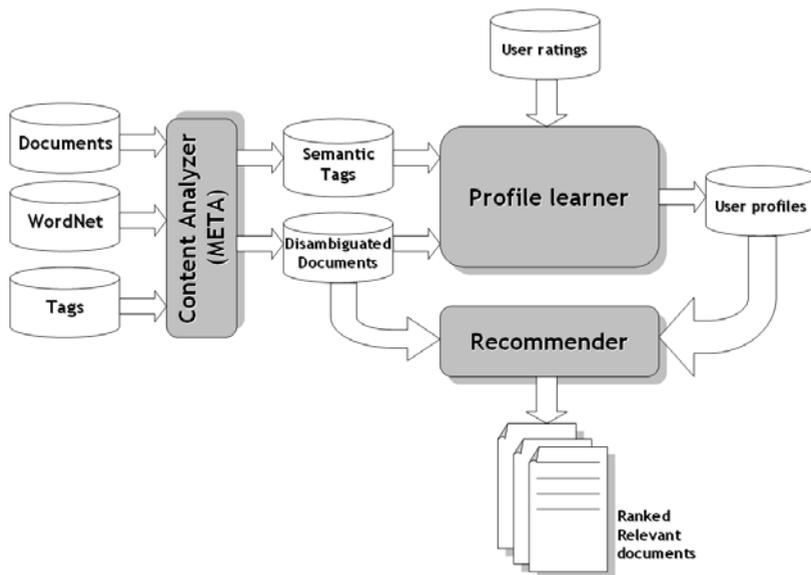


Fig. 2. FIRSt General Architecture

the problems due to the natural language ambiguity. The final outcome of the preprocessing step is a repository of disambiguated documents. This semantic indexing is strongly based on natural language processing techniques, such as Word Sense Disambiguation (WSD), and heavily relies on linguistic knowledge stored in the WORDNET [7] lexical ontology. Semantic indexing of contents is performed by the Content Analyzer, which relies on META (Multi Language Text Analyzer), a natural language processing tool developed at the University of Bari, able to deal with documents in English or Italian [2].

The complete description of the adopted WSD strategy is not described here, because already published in [11].

- PROFILE LEARNER – it implements a supervised learning technique for learning a probabilistic model of user interests from disambiguated documents rated according to her interests. This model represents the semantic profile, which includes those concepts that turn out to be the most indicative of the user preferences. In FIRSt the problem of learning user profiles is casted as a binary Text Categorization task [10] since each document has to be classified as interesting or not with respect to the user preferences. The algorithm for inferring user profiles is naïve Bayes text learning, widely adopted in content-based recommenders. Details about implemented algorithm are provided in [5].
- RECOMMENDER – it exploits the user profile to suggest relevant documents by matching concepts contained in the semantic profile against those contained in documents to be recommended.

The outcome of the experiments conducted in [3] demonstrated that tag integration in the recommendation process causes an increase of predictive accuracy of the recommender. In this work we will continue this analysis, by showing how this higher accuracy can be exploited for personalization goals in a real-world application: artwork recommendation in Vatican Picture Gallery scenario. The strategy adopted to personalize the services for the artwork recommendation scenario might be exploited for the design of recommendation services for e-commerce applications.

4 CHAT Project Overview

CHAT is a research project which aims to developing a platform for multimodal fruition of cultural heritage content.

All the information about a museum is contained in specific data structures and is stored in Fedora digital library¹. The contents in the Fedora are represented by digital objects. In our data modeling (Figure 3) we have three main digital objects: *opus*, *room*, and *author*. Some relations are defined among them:

- *hasCollectionMember* that relates a *room* with an *opus* (the inverse relation is *hasLocation*);
- *isAuthorOf* that relates an *author* with an *opus* (the inverse relation is *hasAuthor*);

Every digital object in Fedora can have one or more datastreams (image, audio, video, text). For each *opus* (painting) we have the following datastreams: *description* (txt or html), *image* (jpg format), *audio* (wav or mp3), *video* (mpeg); *author* and *room* have only textual contents.

All this data are static and not personalized: thus, starting from the same request (for example, a more detailed textual description of an artworks), all users obtain the same answer.

How can we improve the quality of information showed to the visitors?

In the expected scenario every visitor entering the museum is provided with a device (PDA/smart phone) with a specific application installed. Thanks to some localization sensors, it is possible to know in which room of the museum the user is, while coming through the doorway, the visitor can acquire detailed information on each painting in that room.

The core of the CHAT system architecture is the Adaptive Dialog Manager (Figure 4), whose purpose is to manage personalization mechanisms in order to let visitors receive suggestions about artworks they could be interested in.

The Adaptive Dialog Manager embeds many components, called *reasoners*, each of which manages different types of information (about the environment, such as the noise and brightness, about the user, such as the age and interaction speed, and about user tastes) coming from different input channels and localization sensors.

¹ <http://www.fedora-commons.org/>

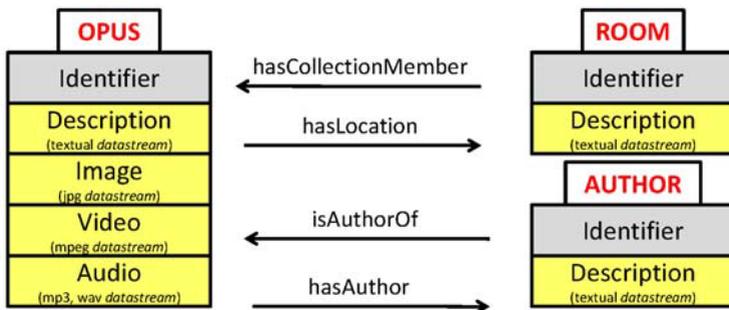


Fig. 3. Data modeling in Fedora Digital Library

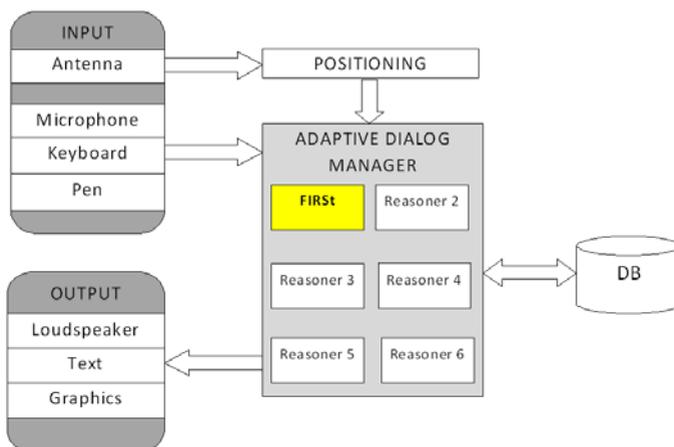


Fig. 4. CHAT Adaptive Dialog Manager for multimodal and personalized fruition

All the data gathered from each reasoner are merged by the Adaptive Dialog Manager, which exploits the *user profile* to find the most interesting items and the most appropriate way to present them, such as just audio or audio and text, video, etc.

4.1 FIRSt@Work in CHAT

FIRSt manages the *content-based profile*, which contains all the information about user preferences on artworks (data that cannot be omitted for personalization goals). The typical steps in a scenario for creating content-based profiles are:

1. Registration to the Museum Portal

In a preliminary phase a user has to subscribe to a dedicated portal. After entering typical demographic data, such as age, sex, instruction, explicit

interests, etc., the user performs a training phase by voting some artworks, randomly chosen from the available ones. After the completion of this step, all user preferences about artworks are stored by FIRSt;

2. Construction of the User Profile

Once finished the training phase, FIRSt builds a profile for each user containing information that turns out to be most indicative of the user preferences;

3. Personalization of Results

When the user visits the museum, he will enjoy additional intelligent personalized services, based on her own user profile.

Specifically, we found two situations where personalized access could improve user experience:

- When data about user interests are gathered, splitting paintings in two disjoint sets, *interesting*, and *not interesting* could be useful to build a personalized tour.
- When the user enters into a room, providing suggestions about paintings she can be interested in can be useful.

To fulfil these specific requirements, we develop two Web Services exposing these functionalities:

- **Tour:** the main idea behind this service is to provide the target user with the subset of items she could be interested in (according to her profile). It is possible to return the first n items for a single user, for example to have the most significant paintings for her. In this service there is not a room-based partition. *This service has been used in CHAT to build a personalized museum tour showing visitors a list of paintings ordered according to their interests.*
- **Room:** it provides the target user with content related to items located in a specific room of the museum. This service takes as input the room identifier (provided by environment sensors) and returns all the artworks in that location.

Room can be adapted to a specific user by exploiting the FIRSt system, that, according to the information stored in the user profile, is able to rank the paintings according to the user interests.

This service has been used in CHAT to discover the most interesting rooms in the museum. This could be very useful when the visitor has no much time to visit the entire museum. In that case, the visitor could plan her visit by starting from the first n locations of interest suggested by the system.

5 Experimental Evaluation

The goal of the experimental evaluation was to understand whether the use of FIRSt in this scenario brings to a substantial improvement of user experience in museum collection fruition. The test has been carried out by using an online

platform which allows registered users to train the system by rating some paintings belonging to the art gallery. The rating scale varies from 1 (dislikes) to 5 (likes). A group of 30 users has been recruited for the test. The task has been divided in three phases:

1. Registration and login;
2. Rating the artworks;
3. Tagging the artworks.

After the registration, each user rated 45 artworks taken from the Vatican Picture Gallery web site and annotated them with a set of free tags. In order to evaluate the effectiveness of the services, we adopted the Normalized Distance-based Performance Measure (NDPM) [16] in order to compare the ranking imposed by the user ratings, with that computed by FIRSt. More specifically, NDPM is used to measure the distance between votes given by a single user u and votes predicted by the system s for a set of items. Given a couple of items t_i, t_j in the Test Set T of an user, the distance between them is calculated through the following schema:

$$\delta_{>u,>s}(t_i, t_j) = \begin{cases} 2 & \iff (t_i >_u t_j \wedge t_j >_s t_i) \vee (t_i >_s t_j \wedge t_j >_u t_i) \\ 1 & \iff (t_i >_s t_j \vee t_j >_s t_i) \wedge t_i \approx_u t_j \\ 0 & \iff \textit{otherwise} \end{cases} \quad (1)$$

The value of NDPM on the Test Set T is calculated through the following equation, where n is the number of couple of items:

$$NDPM_{>u,>s}(T) = \frac{\sum_{i \neq j} \delta_{>u,>s}(t_i, t_j)}{2 \cdot n} \quad (2)$$

For the *Room* service, a single room was set as test set, in order to measure the distance between the ranking imposed on paintings in a room by the user ratings and the ranking predicted by FIRSt.

The methodology of the experiment was as follows:

1. the Training Set (TS_i) for user $u_i, i = 1..30$ is built, by including 50% of all ratings given by u_i (randomly selected);
2. the profile for u_i is built by FIRSt, by exploiting ratings in TS_i ;
3. the profile is used for the computation of the classification scores for the class *likes* for the paintings not included in TS_i ;
4. scores computed by FIRSt and ratings given by users on paintings not included in TS_i are compared.

The test was carried out for 3 rooms in which paintings are located. Generally, NDPM values lower than 0.5 reveal acceptable agreement between the two rankings.

From results reported in Table 1, it can be noticed that the average NDPM is lower than 0.5. In particular, values are lower than 0.5 for 19 users out of 30

Table 1. NDPM for each user (averaged on 3 rooms)

User	NDPM	User	NDPM	User	NDPM	User	NDPM	User	NDPM
u_1	0.56	u_2	0.48	u_3	0.53	u_4	0.65	u_5	0.57
u_6	0.52	u_7	0.38	u_8	0.54	u_9	0.39	u_{10}	0.39
u_{11}	0.46	u_{12}	0.51	u_{13}	0.49	u_{14}	0.46	u_{15}	0.35
u_{16}	0.43	u_{17}	0.45	u_{18}	0.36	u_{19}	0.46	u_{20}	0.35
u_{21}	0.51	u_{22}	0.47	u_{23}	0.39	u_{24}	0.55	u_{25}	0.36
u_{26}	0.46	u_{27}	0.39	u_{28}	0.42	u_{29}	0.60	u_{30}	0.55
								<i>Avg</i>	<i>0.47</i>

(63%), highlighted in bold in the table. Among these users, NDPM for 9 of them is even lower than 0.4, thus revealing that the ranking of paintings proposed by FIRSt is very effective for 30% of the population involved in the test. The main conclusion which can be drawn from the experiment is that the service is capable of providing a quite effective user experience in museum fruition.

6 Related Work

Museums have already recognized the importance of providing visitors with personalized access to artifacts. PEACH (Personal Experience with Active Cultural Heritage) [13] and CHIP (Cultural Heritage Information Personalization) [15] projects are only two examples of the research efforts devoted to support visitors in fulfilling a personalized experience and tour when visiting artwork collections. In particular, the recommender system developed within CHIP aims to provide personalized access to the collections of the Rijksmuseum in Amsterdam. It combines Semantic Web technologies and content-based algorithms for inferring visitors' preference from a set of scored artifacts and then, recommending other artworks and related content topics.

The Steve.museum consortium [14] has begun to explore the use of social tagging and folksonomy in cultural heritage personalization scenarios, to increase audience engagement with museums' collections. Supporting social tagging of artifacts and providing access based on the resulting folksonomy, open museum collections to new interpretations, which reflect visitors' perspectives rather than curators' ones, and help to bridge the gap between the professional language of the curator and the popular language of the museum visitors. Preliminary explorations conducted at the Metropolitan Museum of Art of New York have shown that professional perspectives differ significantly from those of naïve visitors. Hence, if tags are associated to artworks, the resulting folksonomy can be used as a different and valuable source of information to be carefully taken into account when providing recommendations to museum visitors. As in the above mentioned works, we have proposed a solution to the challenging task of identifying user interests from tags. Since the main problem lies in the fact that tags are freely chosen by users and their actual meaning is usually not very

clear, the distinguishing feature of our approach is a strategy for the “semantic” interpretation of tags by means of WORDNET.

7 Conclusions and Future Work

In this paper we investigated about the design of recommendation services based on folksonomies and their concrete exploitation in real-world applications. We evaluated the FIRSt recommender system in the cultural heritage domain, by integrating the system in an adaptive platform for multimodal and personalized access to museum collections. Each visitor, equipped with a mobile terminal, enjoys an intelligent guide service which helps her to find the most interesting artworks according to her profile and contextual information (such as her current location in the museum, noise level, brightness, etc.). Experimental evaluations showed that FIRSt is capable of improving user museum experience, by ranking artworks according to visitor tastes, included in user profiles. The profiles are automatically inferred from both static content describing the artworks and tags chosen by visitors to freely annotate preferred artworks. The personalized ranking allows building services for adaptive museum tours. Since FIRSt is capable of providing recommendations for items in several domains, provided that descriptions of items are available as text documents (e.g. plot summaries, reviews, short abstracts), we will try to investigate its application in different scenarios such as book or movie recommendation.

Acknowledgments

This research was partially funded by MIUR (Ministero dell’Universita’ e della Ricerca) under the contract Legge 297/99, Prot.691 CHAT “Cultural Heritage fruition & e-Learning applications of new Advanced (multimodal) Technologies” (2006-08). The authors are grateful to Massimo Bux for his effort in developing the services and performing the experimental evaluation.

References

1. Balabanovic, M., Shoham, Y.: Fab: Content-based, Collaborative Recommendation. *Communications of the ACM* 40(3), 66–72 (1997)
2. Basile, P., de Gemmis, M., Gentile, A.L., Iaquina, L., Lops, P., Semeraro, G.: META - MultilanguagE Text Analyzer. In: *Proceedings of the Language and Speech Technnology Conference - LangTech 2008*, Rome, Italy, February 28-29, 2008, pp. 137–140 (2008)
3. Basile, P., de Gemmis, M., Lops, P., Semeraro, G., Bux, M., Musto, C., Narducci, F.: FIRSt: a Content-based Recommender System Integrating Tags for Cultural Heritage Personalization. In: Nesi, P., Ng, K., Delgado, J. (eds.) *Proceedings of the 4th International Conference on Automated Solutions for Cross Media Content and Multi-channel Distribution (AXMEDIS 2008) - Workshop Panels and Industrial Applications*, Florence, Italy, November 17-19, 2008, pp. 103–106. Firenze University Press (2008)

4. Burke, R.: Hybrid Recommender Systems: Survey and Experiments. *User Modeling and User-Adapted Interaction* 12(4), 331–370 (2002)
5. de Gemmis, M., Lops, P., Semeraro, G., Basile, P.: Integrating Tags in a Semantic Content-based Recommender. In: *Proceedings of the 2008 ACM Conference on Recommender Systems, RecSys 2008, Lausanne, Switzerland, October 23-25, 2008*, pp. 163–170 (2008)
6. Lops, P., Degemmis, M., Semeraro, G.: Improving Social Filtering Techniques Through WordNet-Based User Profiles. In: Conati, C., McCoy, K., Paliouras, G. (eds.) *UM 2007. LNCS*, vol. 4511, pp. 268–277. Springer, Heidelberg (2007)
7. Miller, G.: WordNet: An On-Line Lexical Database. *International Journal of Lexicography* 3(4) (1990); Special Issue
8. Mladenic, D.: Text-learning and related intelligent agents: a survey. *IEEE Intelligent Systems* 14(4), 44–54 (1999)
9. Resnick, P., Varian, H.: Recommender Systems. *Communications of the ACM* 40(3), 56–58 (1997)
10. Sebastiani, F.: Machine Learning in Automated Text Categorization. *ACM Computing Surveys* 34(1) (2002)
11. Semeraro, G., Degemmis, M., Lops, P., Basile, P.: Combining Learning and Word Sense Disambiguation for Intelligent User Profiling. In: Veloso, M.M. (ed.) *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pp. 2856–2861 (2007) ISBN 978-I-57735-298-3
12. Shardanand, U., Maes, P.: Social Information Filtering: Algorithms for Automating “Word of Mouth”. In: *Proceedings of ACM CHI 1995 Conference on Human Factors in Computing Systems*, vol. 1, pp. 210–217 (1995)
13. Stock, O., Zancanaro, M., Busetta, P., Callaway, C.B., Krüger, A., Kruppa, M., Kuflik, T., Not, E., Rocchi, C.: Adaptive, intelligent presentation of information for the museum visitor in peach. *User Model. User-Adapt. Interact.* 17(3), 257–304 (2007)
14. Trant, J., Wyman, B.: Investigating social tagging and folksonomy in art museums with *steve.museum*. In: *Collaborative Web Tagging Workshop at WWW 2006, Edinburgh, Scotland (May 2006)*
15. Wang, Y., Aroyo, L.M., Stash, N., Rutledge, L.: Interactive user modeling for personalized access to museum collections: The rijksmuseum case study. In: Conati, C., McCoy, K., Paliouras, G. (eds.) *UM 2007. LNCS*, vol. 4511, pp. 385–389. Springer, Heidelberg (2007)
16. Yao, Y.Y.: Measuring retrieval effectiveness based on user preference of documents. *Journal of the American Society for Information Science* 46(2), 133–145 (1995)