

SpIteR: a Module for Recommending Dynamic Personalized Museum Tours

Pierpaolo Basile, Marco de Gemmis, Leo Iaquina,
Pasquale Lops, Cataldo Musto, Fedelucio Narducci, Giovanni Semeraro
Department of Computer Science, University "Aldo Moro"
Bari, Italy

{basilepp, degemmis, iaquina, lops, musto, narducci, semeraro}@di.uniba.it

Abstract—Recommender systems (RSs) proved to make easier the task of accessing relevant information in a broad range of domains. In content-based RSs, preferences on content items expressed by users turned out to be reliable indicators to suggest and filter interesting contents. Item representation plays a key role in content-based RSs, thus choosing proper facets to represent items is a fundamental task for deploying effective RSs. Contextual facets are often marginally relevant to predict user preferences, but in some domains disregarding contextual facets makes recommendations useless.

This paper proposes a strategy to improve the effectiveness of a content-based RS that dynamically suggests tours within a museum by exploiting contextual facets such the physical layout of items and the interaction of users with the environment.

Keywords-Context-aware Recommender Systems; Cultural Heritage Fruition; User-generated Contents

I. INTRODUCTION

The importance of providing digital access to cultural heritage collections has been already acknowledged by museums for almost four decades [1]. More recently, it has been asserted that the personalization drives “the museum monologue” into “a user-centred information dialog” between the museum and its visitors [2]. Visitors spend more and more time to discover interesting artworks, prepare a museum tour, or learn related knowledge about artworks, usually in relation to a (possible) physical museum visit [3]. Cultural heritage personalization refers to supporting visitors in the selection and filtering of preferred artifacts, and in the creation of personalized tours.

Since RSs have proved to be useful in helping users access to desired information (especially in domains where users are not expert or familiar with) [4], they have found their way also in the context of museums. The goal of a RS is finding out the k most interesting items for the current user. Thus, prominent aspects of users and items are scanned and similarity techniques are employed to determine the most relevant suggestions. Among techniques used to design RSs, the most investigated ones concern content or metadata of items in the content-based approach, and the user’s social environment in the collaborative filtering approach.

The user active role with respect to RSs concerns traditionally only to provide a relevance feedback (rating) on items and suggestions. The Web 2.0 (r)evolution is encouraging people to overstep the passive role of content consumers.

Thus, the user increased participation can be exploited in recommendation process. One of the forms of user-generated content (UGC) that has drawn more attention is *tagging*, i.e. resource annotation with free keywords, that builds a socially-constructed classification schema, called *folksonomy* (folks + taxonomy). Social tagging in the cultural heritage personalization scenario promotes the audience engagement with museum collections. Providing access based on the resulting folksonomy opens museum collections to new interpretations, which reflect also visitor perspectives and not only curator ones, and helps to bridge the gap between the professional language of the curator and the popular language of the museum visitor.

Beside encouraging the user increased participation, in order to make user experience enthralling, the item arrangement as well the user interaction with the environment are very relevant. Indeed, a static ordered list in accordance with only the assessed interests is not the best strategy to present the most interesting items, e.g. tortuous paths with repetitive passages make the user disoriented especially under a time constraint. Finally, different users interact with the environment in different manner, e.g. they travel with different speed, they spend different time to admire artworks, they divert from suggested tour.

This paper presents the SpIteR (Spatial Item Recommender) module that suggests dynamic context-sensitive tours according to user profiles learned on static contents and UGCs. Indeed, the module uses contextual facets to adjust (“contextualize”) the resulting set of content-based recommendations in post-filtering [5]. In the “virtuous circle” vision [6], i.e. the connection between on-line and on-site experiences, SpIteR is a preliminary attempt to react to the user behavior contextualizing the personalized tour that is an open issues in cultural heritage fruition, for instance for the CHIP project as recently reported by Wang et al. [3].

The remainder of this paper is structured as follows. Section II describes the our recommender system. Section III discusses the case study. Finally, Section IV draws conclusions and provides direction for future work.

II. ARCHITECTURE

The recommendation process is performed in several steps, each of which is handled by a separate component of

the architecture shown in Figure 1. First, given a collection of documents, a preprocessing step is performed by the *Content Analyzer*, which performs *Word Sense Disambiguation* (WSD) on both static and dynamic content to identify correct senses, corresponding to concepts identified from words in the text. Subsequently, a learning step is performed by the *Profile Learner* on the training set of documents, to generate a sense-based probabilistic model of the user interests. Then, the *Item Recommender* component implements a Naïve Bayes text categorization algorithm to classify documents as interesting or not for a specific user by exploiting the learned probabilistic model. Finally, the *SpIteR* module rearranges the suggested items in a personalized tour using information about the environment and the user behavior.

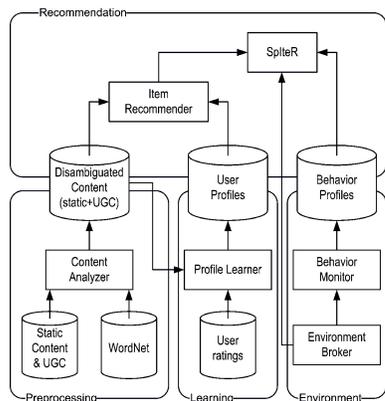


Figure 1. Recommender system architecture

A. Content Analyzer: content semantic indexing

We propose a sense-based representation so that *semantic* user profiles can be learned. The *Content Analyzer* deals with the issue that, although words occur in a text, meanings do not, since they are often hidden in the context. It implements the task of WSD to determine the sense of ambiguous words, adopting WORDNET as sense inventory. The adopted WSD is not described here due to space limitations. More details are reported in [7]. What we would like to point out here is that the WSD allows to obtain a synset-based vector space representation, called bag-of-synsets (BOS), that is an extension of the classical bag-of-words (BOW) model. A BOS vector is obtained for each textual metadata describing the items. For the museum artworks, we use four textual metadata: *artist*, *title*, and *description* slots contain static contents; *tags* slot contains UCGs.

B. Profile Learner: static content and UGC in user profiles

We consider the problem of learning user profiles as a binary Text Categorization task since each document has to be classified as interesting or not with respect to the user preferences. Therefore, the set of categories is restricted to c_+ , that represents the positive class (user-likes), and c_- the negative one (user-dislikes). The induced probabilistic model is used to estimate the *a posteriori* probability,

$P(c_j|d_i)$, of document d_i belonging to class c_j . The algorithm adopted for inferring user profiles follows a Naïve Bayes text learning approach, widely used in content-based recommenders [8], which is not described here due to space limitations. The profile contains the user identifier and the *a-priori* probabilities of liking or disliking an item, apart from its content. Moreover, the profile is structured in two main parts: *profile_like* contains features describing the concepts able to deem items relevant, while features in *profile_dislike* should help in filtering out not relevant items. Each part of the profile is structured in slots, resembling the same representation strategy adopted for artwork metadata. Each slot reports the features (WORDNET identifiers) occurring in the training examples, with corresponding frequencies computed in the training step. Frequencies are used by the Bayesian learning algorithm to induce the classification model (i.e. the user profile) exploited to suggest relevant artworks in the recommendation phase.

C. SpIteR: tour model and strategies

The tour suggestion task requires also knowledge about the item layout. We propose to basically represent items as nodes of an Euclidean graph. Thus the museum tour is quite similar to the classical *Traveling Salesperson Problem* (TSP). TSP is a well known combinatorial optimization problem and it has been studied extensively in many variants. The problem is NP-hard, as shown by Papadimitriou [9], even if weight of each edge satisfies the triangle inequality (in addition to nonnegative and symmetry properties) like in the Euclidean graphs. Further information about museum tour items can be exploited to obtain the problem solution. Usually, few items are placed in rooms and each room is connected with some other rooms. Rooms provide a simpler perspective from combinatorial complexity point of view, but, while tour visits each item at most once (like in k -TPS), each room can be visited more times. Thus, a *divide et impera* strategy is suitable: a first step to search sequences of rooms with recommended items, a second step to locally solve TSP for each room, and finally a composition step to achieve sequences of recommended items.

The user profile allows to reduce the involved nodes. Indeed, the suggested tour should consist of recommended items, i.e., the k most interesting items for the current user. The question about k value selection arises: Intuitively, it depends on how long should be the personalized tour, e.g., the preferred tour duration and the user behavior must be counted. Thus the estimate of k value mainly deals with overall time constraint on expected time spent by the user to admire the supposed most interesting museum items. That estimate is allowed to be rough (i.e., the actual time spent to reach items can be disregarded) since the early goal is to cut back the supposed less interesting nodes. Speed and stay times are parameters related to user behavior. At the beginning, they are estimated on the basis of stereotypical

user profile [10] and then updated according to data collected during the tour from the actual user behavior. Hence the k value can change during the tour: that is not an issue, since the tour is constantly monitored to detect significant deviations from proposed one. When the user behavior requires too many updates to the behavior profile or the user skips an item or she stays in front of un-recommended items, the tour is planned again taking into account the previous user behavior and the actual viewed items.

The actualized k value allows to discover rooms with further recommended items and therefore *blind alleys* can be pruned ever since the first step concerning the room sequences. Thus the actual distribution of the k most interesting items among rooms reduces the search space of the first step as well as of the second one. Indeed, the local search within each room is focused only on items that are not still suggested or viewed.

Also the composition step to achieve sequences of recommended items shows the effects of the focus on the k more interesting items. Indeed, there are cuts on both room sequences and intra-room item sequences. Therefore the composition refers to a simplified version of the initial problem that leads in few seconds to the optimal tour covering the most interesting items according to the user profile.

III. EXPERIMENTAL EVALUATION

The goal of the experimental evaluation is twofold: to comparing the predictive accuracy of provided recommendations and to evaluate the context effects on tours.

Thus, user profiles are learned from static content only as well as from both static content and UGC. In addition, to properly investigate the effects of including social tagging in the recommendation process, a distinction has to be made between considering, for an artifact rated as interesting by a user, either the whole folksonomy (i.e., the community tags used by all visitors to annotate that artifact), or only the tags entered by that user for that artifact (i.e., the user's contribution to the whole artifact folksonomy).

The evaluation of the context effects on tours is mainly focused on time percentage spent to move. Indeed the time constraint is the ultimate reason to adopt a RS in the museum scenario. We assume that artwork admiring time is not related to any order in the artwork sequence and, therefore, the item rearrangement makes sense to optimize the tour.

The dataset consists of 45 paintings chosen from the collection of the Vatican picture-gallery¹. All artworks are laid in a schematic environment model to deal with the spatial layout influence on recommending. We involve 30 users who volunteered took part in the experiments. The average age of the users is in the middle of twenties. None of the users is an art critic or expert.

Users are requested to express their preferences and to annotate the paintings: 4300 tags are used.

Since the content-based recommending component is conceived as a text classifier, its effectiveness is evaluated by *Precision*, *Recall* and *F1 measure*. The methodology adopted for obtaining training and test sets is the K -fold cross validation, with $K = 5$. Different content types are used:

- EXP #1: STATIC CONTENT - only title, artist and description of the painting
- EXP #2: PERSONAL TAGS - only tags provided by a specific user on a specific painting
- EXP #3: SOCIAL TAGS - all the tags provided by all the users on a specific painting
- EXP #4: STATIC CONTENT + PERSONAL TAGS
- EXP #5: STATIC CONTENT + SOCIAL TAGS

In the evaluation of the context effects on tours, the time constraint brings about the number of recommended items. A very short time (i.e., a constraint selecting very few items) as well as a very long time (i.e., a constraint selecting almost all items) are not particularly suitable to state layout influence. Thus, the time constraints are established so that roughly 10, 20 and 30 items have to be arranged in personalized tours. The comparison is performed among the suggested tour path and the path resulting from the ranked list of the same items as stated by the content-based user profile. Figure 2 and Figure 3 show two samples of paths with the same items².

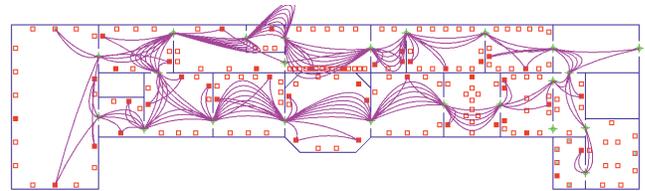


Figure 2. Sample of path resulting from the ranked suggestion list from the content-based user profile

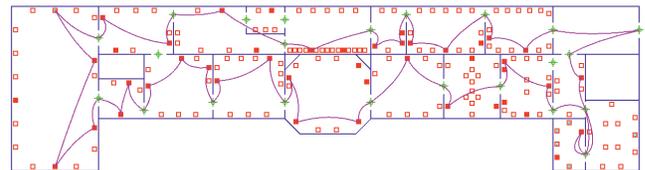


Figure 3. Sample of actually recommended tour path

Table I shows that the integration of UGC (whether social or personal tags) causes an increase of precision in the process of recommending artifacts to users. More specifically, precision of profiles learned from both static content and tags (hereafter, augmented profiles) outperformed the precision of profiles learned from either static content (hereafter, content-based profiles) or just tags (hereafter, tag-based profiles). The improvement ranges between 2% and 2.40%. The comparison of the precision of content-based profiles with

¹http://mv.vatican.va/3_EN/pages/PIN/PIN_Main.html

²bend lines are only a graphical expedient to avoid hidden segments

Table I
RESULTS OF THE K -FOLD CROSS VALIDATION

	Type of content			Precision	Recall	F1	10 Items	20 Items	30 Items
	Static	Personal Tags	Social Tags						
Exp #1	X			75.86	94.27	84.28	-48.1%	-62.9%	-74.3%
Exp #2		X		75.96	92.65	83.26	-45.9%	-58.5%	-71.2%
Exp #3			X	75.59	90.50	82.13	-49.7%	-63.7%	-73.7%
Exp #4	X	X		78.04	93.60	84.93	-47.4%	-62.9%	-73.7%
Exp #5	X		X	78.01	93.19	84.73	-50.3%	-64.9%	-74.3%

that of tag-based profiles suggests that just tags are sufficient for providing accurate recommendations, while a decrease of recall (-1.62% with personal tags, -3.77% with social tags) actually shows that static content cannot be neglected even if tags are available. The higher decrease of recall registered with social tags leads to conclude that community tags introduce some noise in the recommendation process (relevant paintings are filtered out due to wrong advice by other users). The general conclusion of the comparison between content-based profiles and augmented profiles is that a significant increase of precision corresponds to a slight and physiological loss of recall. The overall accuracy of augmented profiles (F1 almost 85%) is considered satisfactory.

Comparing estimated times spent on the two type of path, the experiments confirm the intuition that item layout allows to reduce the moving time rate over the total tour time. This reduction is more significant when the number of filtered items increases, but it shows to be marginally sensitive to the type of content used to build the user profiles, therefore the tour optimization is not decisive on choosing the content-based approach.

IV. CONCLUSIONS AND FUTURE WORK

In this paper we have investigated how to effectively combine existing content-based filtering algorithms with UGC and how to optimize the recommendations according to context aspects in the cultural heritage personalization.

The main contribution of the paper is a strategy to enable a content-based RS to use also the physical layout of items and the user interaction with the environment. In particular, the developed content-based RS utilizes machine learning techniques to infer user profiles both from static content and UGC. The carried out experiments highlight that the integration of UGC causes an increase of precision in the process of recommending artifacts to users and that the optimization on moving time is more sensitive to the number of filtered items than to the content type. However, the findings from the experiments with social tags run counter our expectation because, as compared to the use of personal tags only, a decrease of precision and recall is observed. To gain more insights on the effects of community-generated content, we need to 1) perform an analysis of what tags are used to build the folksonomies and how they affect the

user profile generation; 2) replicate the experiments with a more heterogeneous community, involving experts in the art domain so as to identify differences with the tagging activity of naïve users. A more heterogeneous community allows also to gain more insights on the context effects, especially about user behavior.

REFERENCES

- [1] E. Ellin, "Museums and the computer: An appraisal of new potentials," *Language Resources and Evaluation*, vol. 4, no. 1, pp. 25–30, 1969.
- [2] J. P. Bowen and S. Filippini-Fantoni, "Personalization and the web from a museum perspective," in *Museums and the Web*, 2004, pp. 63–78.
- [3] Y. Wang, L. Aroyo, N. Stash, R. Sambeek, Y. Schuurmans, G. Schreiber, and P. Gorgels, "Cultivating personalized museum tours online and on-site," *Interdisciplinary Science Reviews*, vol. 34, no. 2, pp. 141–156, 2009.
- [4] P. Resnick and H. R. Varian, "Recommender systems," *Communications of the ACM*, vol. 40, no. 3, pp. 56–58, 1997.
- [5] G. Adomavicius and A. Tuzhilin, "Context-aware recommender systems," in *RecSys'08: Proc. of the ACM Conference on Recommender Systems*, 2008, pp. 335–336.
- [6] A. Barry, "Creating a virtuous circle between a museums on-line and physical spaces," in *Museums and the Web 2006*, 2006. [Online]. Available: <http://www.archimuse.com/mw2006/papers/barry/barry.html>
- [7] G. Semeraro, M. Degemmis, P. Lops, and P. Basile, "Combining learning and word sense disambiguation for intelligent user profiling," in *Proc. of the 20th Int. Joint Conference on Artificial Intelligence IJCAI-07*, 2007, pp. 2856–2861.
- [8] D. Mladenic, "Text-learning and related intelligent agents: a survey," *IEEE Intelligent Systems*, vol. 14, no. 4, pp. 44–54, 1999.
- [9] C. H. Papadimitriou, "The euclidean traveling salesman problem is np-complete," *Theor. Comput. Sci.*, vol. 4, no. 3, pp. 237–244, 1977.
- [10] B. Shapira, P. Shoval, and U. Hanani, "Stereotypes in information filtering systems," *Information Processing and Management*, vol. 33, no. 3, pp. 273–287, 1997.