Exploiting Image Segmentation Techniques for Social Filtering of Educational Content

Pythagoras Karampiperis[1] and Aristeidis Diplaros[2]

[1] Department of Technology Education and Digital Systems,
University of Piraeus, Greece,
pythk@ieee.org
[2] Informatics Institute, University of Amsterdam, The Netherlands
diplaros@gmail.com

Abstract. The need for applying advanced social information retrieval techniques for personalizing web-based information discovery has been identified as a key challenge. Until now, significant R&D effort has been devoted aiming towards applying collaborative filtering techniques for educational content retrieval. However, limited attention has been given to the use of educational metadata as a mean to enhance social filtering techniques via educationally informed filtering decisions. In this paper we propose the use of an add-on filtering service on existing social filtering systems/applications so as to create a data post-filtering mechanism that makes use of intelligence stored in TEL metadata. The proposed methodology starts with the generation of a matrix that represents the educational characteristics of the resources suggested by typical social filtering techniques and applies post-filtering using the educational "footprint" of the resources already used by the targeted end-user.

Keywords: Technology Enhanced Learning, Educational Metadata, Social Filtering, Data Clustering.

1 Introduction

The high rate of evolution of Web 2.0 applications implies that on the one hand, increasingly complex and dynamic web-based learning infrastructures need to be managed more efficiently, and on the other hand, new type of learning services and mechanisms need to be developed and provided. To meet the current needs, such services should satisfy a diverse range of requirements, as for example, personalization based on social filtering [1].

In this context, the need for applying advanced social information retrieval techniques for personalizing web-based information discovery and retrieval has been identified as a key challenge. This has become more critical in the case of Technology Enhanced Learning applications, since on the Web a vast variety of digital learning resources exist that have the potential to facilitate teaching and learning tasks. Until now, significant R&D effort has been devoted aiming towards applying collaborative filtering techniques for educational content retrieval [2]. These techniques are using usage log files over a set of educational resources to provide personalized recommendations by comparing the profile of the learner in hand with similar persons/groups recorded in the historical log data [3, 4, 5]. However, limited attention has been given to the use of educational metadata as a mean to enhance social filtering techniques via educationally informed filtering decisions.

In this paper we propose the use of an add-on filtering service on existing social filtering systems/applications so as to create a data post-filtering mechanism that makes use of intelligence stored in TEL metadata. The main driver of this work was inspired by the idea of using visualization information for accessing Learning Object Repositories [6]. Our goal was to investigate how image segmentation techniques could be applied in order to enhance the social filtering process of educational content. More precisely, the proposed methodology starts with the generation of a matrix that represents (in visual form) the educational characteristics of the resources suggested by typical social filtering techniques and applies post-filtering using the educational "footprint" of the resources already used by the targeted end-user. For the generation of the resource filter we utilize image segmentation techniques, taking into account the spatial coherence of the created visual representation. We treat the filtering problem as an inference problem, assuming that each pixel in the educational "footprint" (visualization) has a hidden binary label associated with it which specifies if it is appropriate for the targeted learner or not. In order to solve the inference problem, we use a variation of the EM algorithm [7] which incorporates the spatial constraints with just a small computational overhead [8].

Moreover, a potential drawback when applying social filtering techniques is that the models used are not fully transparent to the end user, thus, affecting the end-users' trust on the provided recommendations [9]. Since the generated filter by the proposed approach is represented visually, end-users can directly observe the core of the educational filtering process and make modifications/updates if desired.

The paper is structured as follows: In section 2, we discuss how educational metadata could be used in order to generate the educational "footprint" (visualization) of a set of educational resources. Section 3 presents the proposed methodology for generating the post-filter for the resources recommended by typical social filtering techniques, using as an input the educational "footprint" of the resources already used by the targeted end-user. Finally, we demonstrate the application of the proposed visualization and filtering process on an easy-to-understand real life scenario.

2 Social Filtering via Educational Metadata Visualizations

Social filtering is a method for making automatic predictions (filtering) about the preferences of a user by collecting preference information from many users. The underlying assumption of social filtering is that the users with similar preferences in the past tend to have similar preferences in the future. There exist three main types of social filtering: active filtering, passive filtering and item-based filtering. Active filtering uses a peer-to-peer approach, based on explicit user ratings over a set of available digital resources. On the other hand, passive filtering uses preference information that was implicitly collected via usage log files. Implicit filtering relies on the historical actions of users to determine a value rating for digital content. Finally, in the case of item-based filtering, items (digital resources) are rated and used as parameters instead of users. This type of filtering uses the ratings to group various items so as to enable potential users to compare them.

Our proposed method is an add-on filtering service on existing passive social filtering systems/applications, utilizing intelligence stored in TEL metadata. The main idea of the proposed approach is to post-filter the recommendations provided by typical passive social filtering techniques using the educational "footprint" of the resources already used by the targeted end-user. To achieve this, we create a matrix that represents (in visual form) the educational characteristics of the resources already recorded in the historical log files. Based on this matrix, we generate another matrix that represents the educational preferences of the targeted user. The latter matrix acts as an educational post-filter on the resources suggested by a typical social filtering system. This post-filtering is made by comparing the generated filter with the educational "footprint" of the resources suggested by a passive social filtering technique. Next paragraphs present how educational metadata are used to create the educational "footprint" of a single resource, as well as, of a set of resources. It is clear that this method is used for creating both the educational representation of the resources already used by the targeted user (which is the input for the filtering generation process), and the educational representation of the resources suggested by a passive social filtering technique (which is the input for the post-filtering process).

2.1 Creating the Educational Footprint of a Learning Resource

In order to generate the educational footprint (representation) of an educational resource we use the corresponding metadata record, a subset of the IEEE Learning Object Metadata (LOM) standard elements. The metadata elements used were selected in such a way that each element uses a specific state vocabulary, as illustrated in Table 1.

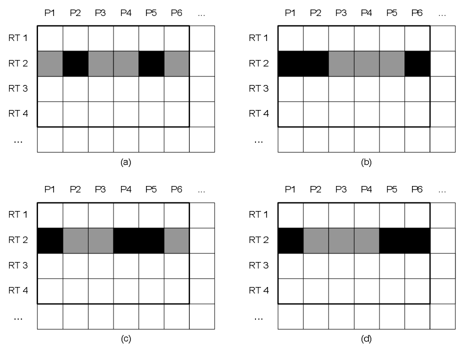

Fig. 1. Examples of representing the educational footprint of individual learning resources with Learning Resource Type (RT) equal to "simulation".

The educational footprint of a learning resource is a 15x8 pixels image where the first dimension (lines) stands for the states of the Learning Resource Type attribute and the second dimension (columns) stands for the rest eight attributes used. Each pixel is colored according to the value of the corresponding attribute of the second dimension. The color coding used for each metadata attribute j is defined by the formula:

$$Color_{RED}^{j} = Color_{GREEN}^{j} = Color_{BLUE}^{j} = \left(1 - \frac{k^{j}}{N}\right) \times 255,$$

where N stands for the number of vocabulary states of metadata attribute j, and k^{j} is the state code of attribute j for a given educational resource.

Table 1. Educational Resource Description Model and Color Coding used.

Metadata Element Used	Vocabulary State	State Code	Color Code (R-G-B)=(X-X-X)	Color
Interactivity Type	active	1	X=(2/3)*255	
	expositive	2	X=(1/3)*255	
	mixed	3	X=0	
Interactivity Level	very low	1	X=(4/5)*255	
	low	2	X=(3/5)*255	
	medium	3	X=(2/5)*255	
	high	4	X=(1/5)*255	
	very high	5	X=0	
Semantic Density	Same Vocabulary and Color Coding with "Interactivity Level"			
Typical Age Range	K12	1	Custom Vocabulary (not defined in IEEE LOM). In our simulations we used the same Color Coding with "Interactivity Type"	
	13-18	2		
	Adults	3		
Difficulty	Same Vocabulary and Color Coding with "Interactivity Level"			
Intended End User Role	teacher	1	X=(3/4)*255	
	author	2	X=(2/4)*255	
	learner	3	X=(1/4)*255	
	manager	4	X=0	
Context	school	1	Same Color Coding with "Intended End User Role"	
	higher education	2		
	training	3		
	other	4		
Typical Learning Time	Custom Vocabulary (not defined in IEEE LOM). In our simulations we used the same Vocabulary and Color Coding with "Interactivity Level"			
Learning Resource Type	exercise	1	This metadata element was used as the second dimension for the creation of the resource visual matrix. Thus, no color coding was used for this metadata element since each line (or set of lines) in the visual matrix represents directly the value of the "Learning Resource Type"	
	simulation	2		
	questionnaire	3		
	diagram	4		
	figure	5		
	graph	6		
	index	7		
	slide	8		
	table	9		
	narrative text	10		
	exam	11		
	experiment	12		
	problem statement	13		
	self assessment	14		
	lecture	15		

Fig.1, presents examples of the produced representations for different cases of educational content, with the same learning resource type. For presentation simplicity, we have used resources that use only two values (states) per each metadata attribute (represented with gray and black colors accordingly).

2.2 Creating the Educational Footprint of a Set of Learning Resources

In order to generate the representation of a set of learning resources, we start from the representation of the first learning resource in the set and extend the resolution of the generated image for each $n \times n$ resources, with $n \geq 2, n \in N^*$ per learning resource type. So the size of the generated representation for a set can be: $(15k) \times (8k)$ pixels, where $k \in N^*$. As a result the generated visualizations can be (15 x 8), (30 x 16), (45 x 24), … pixels. Fig.2, presents the aggregated representation of the resources demonstrated in previous section (Fig.1).

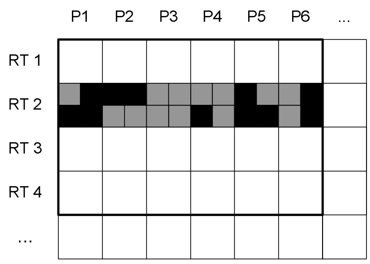

Fig. 2. Example of aggregated representation of a set of learning resources.

Next section presents the methodology for generating the educational post-filter (that is, a matrix which represents the educational preferences of the targeted user) for the resources suggested by a typical passive social filtering system.

3 Generating the Filter for Educational Resource Post-Filtering

The core idea of the filtering generation method used in this paper, is to treat the pixel labels of a representation as independent random variables from a common prior distribution p(si) (which we are going to learn by the EM algorithm), but constrain their posterior distributions (computed in the E-step of the EM algorithm) according to the spatial dependencies between pixels [8]. Although educational metadata properties are correlated, the idea of treating them as independent random variables seems (from preliminary investigation) that it does not affect the filtering process. Of course, this issue will be a subject for deeper investigation in the future, since in this paper our goal was to setup the framework for educational post-filtering of social filtering processes rather than the deep comparison of data clustering techniques to handle the correlation of educational metadata.

In particular, we define a log-likelihood function:

$$L(\theta) = \sum_{i=1}^{n} \log \sum_{s_i} p(c_i \mid s_i) p(s_i) \quad (1)$$

where the parameter θ summarizes all unknown parameters in the model. These unknown parameters are learned by the EM algorithm [10]. More precisely, θ includes the prior probability of each state of the educational metadata parameters. In order to capture the spatial constraints of the pixel labels into an EM algorithm, we employ a variational approximation in which we maximize in each step a lower bound of $L(\theta)$. This bound $F(\theta, Q)$ is a function of the current mixture parameters θ and a factorized distribution $Q = \prod_{i=1}^{n} q_i(s_i)$, where each $q_i(s_i)$ corresponds to pixel i but defines an otherwise arbitrary discrete distribution over s_i.

An attractive property of the variational EM framework is that in each step of the algorithm we are allowed to assign any distribution $q_i(s_i)$ to individual pixels as long as this increases the energy F. In summary, our variational EM algorithm is as follows:

1. (Initialization) Start with a random guess for the parameter vector θ.
2. (Standard E-step) Compute the Bayes posterior probabilities over pixel labels given the pixel colors given the current estimate of θ.
3. Smooth the responsibilities of neighboring pixels by applying a local filter on the set of assigned posteriors (and then renormalize if needed). An efficient way to do this is to represent the set of assigned responsibilities as an image and apply a standard Gaussian smoothing filter.
4. (Standard M-step) Use the smoothed responsibilities in order to update the parameter θ as in standard EM [9]. If convergence stop, else go to step 2.

4 Demonstration

In order to make a preliminary evaluation of the effectiveness the proposed approach we used 10 Learning Object sets consisting of 135 learning object metadata records, that is, 9 Learning Objects per Learning Resource Type (simulating 10 different end-user's historical log files) and a set of 20 learning object metadata records (simulating recommendations from a passive social filtering system), with normal distribution over the value space of each metadata element. The goal of the evaluation was to test the ability of filtering out learning resources with educational footprint that does not match the educational preferences of a given end-user. From this preliminary evaluation, we have evidence that such an add-on service has the potential to enhance social filtering techniques via educationally informed filtering decisions.

Fig.3 presents an example of how the educational footprint for a set of 9 Learning Objects per Learning Resource Type is generated, depicting the step-by-step result of this process for the case of "Interactivity Type" metadata attribute. As we can observe, this is an incremental process starting with the representation of the educational footprint of the first learning object in the set (Fig.3a), continues with the representation for the first 2x2 learning objects (Fig.3b), then with the representation of the first 3x3 learning objects (Fig.3c), and so on for larger sets of learning objects (Fig.3d).

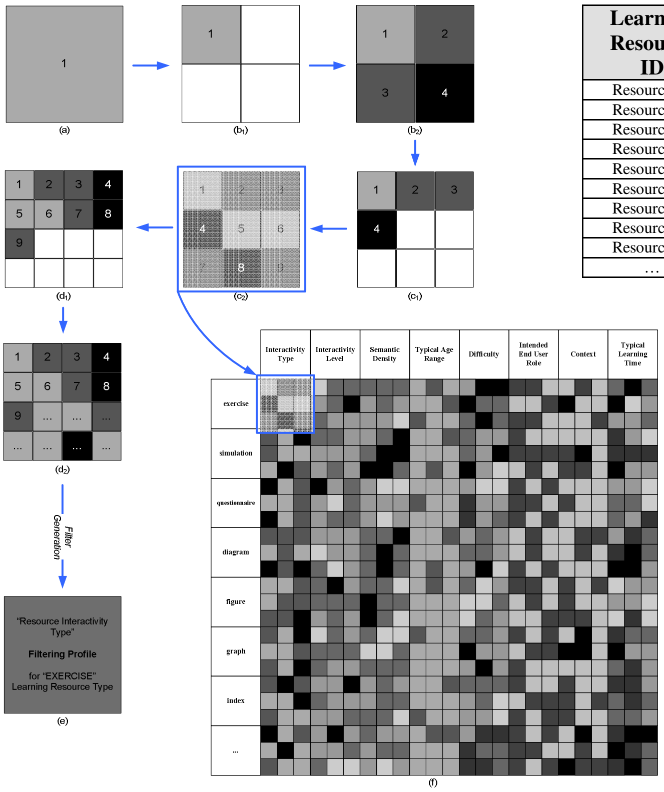

Learning Resource ID	Resource Interactivity Type	Color Code
Resource #1	active	
Resource #2	expositive	
Resource #3	expositive	
Resource #4	mixed	
Resource #5	active	
Resource #6	active	
Resource #7	expositive	
Resource #8	mixed	
Resource #9	expositive	
…	…	…

Fig. 3. Generating the educational footprint of a set of learning resources.

This representation is used as an input for generating the resource filter for the educational post-filtering of the resources suggested by a typical social filtering system/application. An example of such a filter is presented in Fig.4.

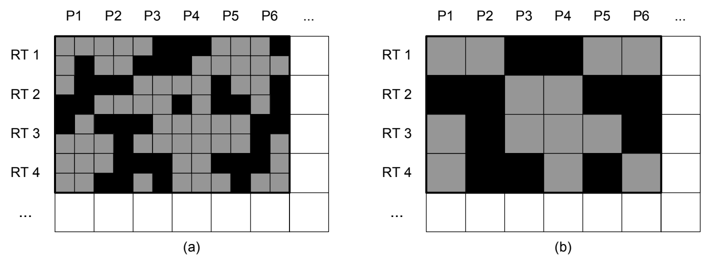

Fig. 4. (a) Example of representing a set of 16 learning objects – 4 per each learning resource type, (b) result of the proposed algorithm acting as a post-filter for future recommendations.

5 Conclusions

In this paper we propose the use of an add-on filtering service on existing social filtering systems/applications so as to create a data post-filtering mechanism that makes use of intelligence stored in TEL metadata. The main driver of this work was inspired by the idea of using visualization information for accessing Learning Object Repositories. Our goal was to investigate how image segmentation techniques could be applied in order to enhance the social filtering process of educational content. The proposed methodology starts with the generation of a matrix that represents the educational characteristics of the resources suggested by typical social filtering techniques and applies post-filtering using the educational "footprint" of the resources already used by the targeted end-user. We treat the filtering problem as an inference problem, assuming that each pixel in the educational content visualization has a hidden binary label associated with it which specifies if it is appropriate for the targeted learner or not. In order to solve the inference problem, we use a variation of the EM algorithm which incorporates the spatial constraints with just a small computational overhead.

References

1. Ahn, J.-W., Farzan, R., and Brusilovsky, P.: Social Search in the Context of Social Navigation. Journal of the Korean Society for Information Management, vol. 23 (2), pp. 147--165 (2006)
2. Recker, M. M., Walker, A., and Wiley, D. A.: Collaboratively filtering learning objects. In D. A. Wiley (Ed.), The Instructional Use of Learning Objects: Online Version (2000)
3. M. Craven, D. DiPasquo, D. Freitag, A. McCallum, T. Mitchell, K. Nigam, and S. Slattery.: Learning to construct knowledge bases from the world wide web. Artificial Intelligence, vol. 118(1-2), pp. 69--113 (2000)
4. B. Mobasher, R. Cooley, and J. Srivastava.: Automatic personalization based on web usage mining. Communications of the ACM, vol. 43(8), pp. 142--151 (2000)
5. B. Mobasher, H. Dai, and M. Nakagawa T. Luo.: Discovery and evaluation of aggregate usage profiles for web personalization. Data Mining and Knowledge Discovery, vol. 6, pp. 61--82 (2002)
6. Klerkx, J., Duval, E., Meire, M.: Using Information Visualization for Accessing Learning Object Repositories. In Proc. of the 8th IEEE International Conference on Information Visualisation, pp. 465 -- 470 (2004)
7. A. P. Dempster, N. M. Laird and D. B. Rubin.: Maximum likelihood from incomplete data via the EM algorithm, J. Roy. Statist. Soc. B, vol. 39, pp. 1--38 (1977)
8. A. Diplaros, N. Vlassis, and T. Gevers.: A spatially constrained generative model and an EM algorithm for image segmentation. IEEE Transactions on Neural Networks, vol. 18(3), pp. 798--808 (2007)
9. D. Pierrakos, G. Paliouras, C. Papatheodorou, and C.D. Spyropoulos.: Web usage mining as a tool for personalization: A survey. User Modeling and User-Adapted Interaction, vol. 13, pp. 311--372 (2003)
10. R. M. Neal and G. E. Hinton.: A view of the EM algorithm that justices incremental, sparse, and other variants, .Learning in graphical models, M. I. Jordan, Ed. Kluwer Academic Publishers, pp. 355--368 (1998)