What if annotations were reusable: a preliminary discussion

Nikos Manouselis[1], Riina Vuorikari[2]

[1] Greek Research & Technology Network
56 Messogion Str.
Athens, Greece
{nikosm}@ieee.org
[2] European Schoolnet, Rue de Treves, 61,
1040 Brussels, Belgium
{Riina Vuorikari}@eun.org

Abstract. This paper discusses the rationale for the representation of user feedback in a structured and reusable format so that it can be reused by different recommender systems. We emphasize how information about the context can be included in such a representation. This work-in-progress takes place in the context of two large European initiatives that set up collections of digital educational resources in distributed repositories to serve the needs of different user communities, and to collect user feedback such as ratings, bookmarks and tags related to the resources. The overall aim is to facilitate the exchange and reuse of their data sets in order to support recommendation of appropriate resources to the end users.

Keywords: Learning resource metadata, annotations, interoperability, reusability, data set.

1 INTRODUCTION

A representation of implicit or explicit feedback from the users regarding the candidate items is required by a recommender system to produce a recommendation. This feedback can be in several forms. For example, in the case of collaborative filtering systems it can be ratings or votes (i.e. if an item has been purchased, viewed or bookmarked). In the case of content-based recommenders, it can be product reviews or simple tags (keywords) that users provide for items. Additional information is also required such a unique way to identify who provides this feedback (user ID) and upon which item (item ID). The user-rating matrix [1] used in collaborative filtering systems is a well-known example.

User feedback representations are stored as data sets that can be used in a number of ways. A very popular use is in the context of evaluation experiments. Published data sets, such as the MovieLeans and EachMovie ones, are very often used in experiment testing the performance of new algorithms [2]. Less often, such

experiments are based on usage data sets from particular applications (e.g. to support the needs of an existing e-commerce site). This means that in many evaluation experiments, testing takes place using data sets that are different than the ones of the potential application context. We believe that a structured way to represent the different types of user feedback in a reusable and interoperable format, while also maintaining information about the context in which this user feedback has been collected, could prove of particular value.

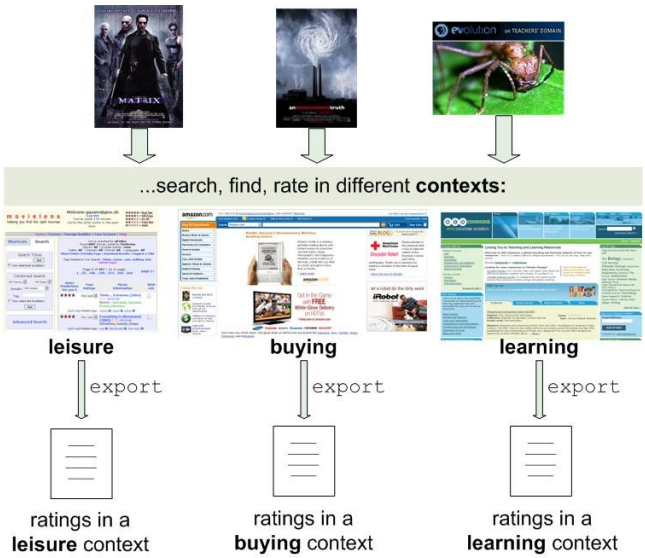

Fig. 1. Ratings on the same items from different contexts.

This would be particularly useful in scenarios where different online environment want to exchange the feedback that users have provided upon the same items, so that they can enhance the available data for their recommender systems. We deal with such situation in two European initiatives that both collect user feedback on digital educational resources: the Metadata Ecology for Learning and Teaching (MELT, http://info.melt-project.eu) and the Organic.Edunet (http://www.organic-edunet.eu) eContentplus projects. In this paper we describe our first steps towards deciding upon a common representation format so that the two projects can collect, store, and exchange data sets containing ratings, bookmarks and tags by their users.

2 PROBLEM DESCRIPTION

The main questions discussed in this paper are: (a) can we find a way to represent and store data sets with user feedback in a structured, interoperable and reusable format?; and (b) is it safe to assume that some given data set with user feedback, which is stored in such a structured, interoperable and reusable way, can be used in other context, such as a new recommender system of a different application context?

Figure 1 illustrates one of the main problems to be addressed in a rather simplistic way. In this figure, the same movie can be found as an item (e.g. for a review, purchase, viewing, downloading) in three different application contexts: one movie

recommender system, one e-commerce site, and one educational portal. Users in all three sites are rating the same movies, depending on the potential use they envisaged. For example, the user of the movie recommender rates the movie according to how much he appreciated/enjoyed it, independently on whether this movie can be easily found or its price. On the other hand, using an e-commerce site a user may rate the movie considering the availability and cost of the particular product package (e.g. DVD). Moreover, on the educational portal, the movie will be rated according to its potential value as a teaching or learning aid.

This simple example illustrates why we believe user-provided feedback can be of particular value when made available for reuse from other systems, but after considering that the two systems may be collecting feedback in a totally different context. For instance, this example describes why it would be probably misleading to use the rating data captured by the movie recommender to test the performance of a recommender system in an educational context. On the other hand, one cannot overlook the value that an existing user feedback (e.g. past ratings) has, as they are upon items that are candidate for recommendation in both systems. To this end, we argue that (ideally) user feedback:

(a) should be represented in general in a structured, reusable and interoperable format that can be shared among different systems;

(b) should incorporate some information about the context in which it has been collected.

3 CAPTURING USER FEEDBACK

Capturing and taking advantage of users' actions on the Web has come a long way since business models were first implemented around the idea of click-stream in the '90. Instead of having the commercial sites taking advantage of the attention that users pay to different products, in the recent years the tide has turned arguing that interactions with the content (e.g. buying, listening, reading feeds) and users reactions to that content (e.g. ratings, reviews, tags) should be something that the user can control. AttentionTrust.org (http://www.attentiontrust.org/), for example, calls this "attention data" and argues that it is a valuable resource that reflects user's interests, activities and values, thus serves as a proxy for their attention.

The first elaboration of AttentionXML for the technology enhanced learning domain has been the Contextualized Attention Metadata (CAM) schema [4]. The CAM schema allows capturing observations about users activities in any kind of tool (not just a browser or newsreader as AttentionXML), with a particular focus on educational software. It has been initially designed to store information about what has attracted users' attention when working with several tools, and to store the interactions of the user with these tools. It also allows storing some basic information about comments, tags or ratings that a user provides when viewing a particular item.

Figure 2 presents an excerpt of XML code storing information according to CAM. It illustrates how ratings and tags are stored within the schema. More specifically, when a user interaction takes place, an event is recorded. This event belongs to a particular actionType, such as 'evaluate' (when a rating is given) or 'tag'. The information that the user provides is stored as an entry with some name and content values. It may also be compliant to some describing schema (e.g. if a particular rating scheme is followed). The figure presents only an example of CAM record, which could also store additional information about the performed action and its input (e.g. the language of a tag).

Attention metadata such as the one stored by CAM can be used for recommendation purposes, since it includes all types of user feedback (e.g. ratings,

bookmarks, purchases, tags, reviews). On the other hand, in order to collect usage data that has been stored with some specification such as CAM, the following requires attention:

(a) To define the exact type and structure of the information stored (e.g. if it is a rating value, what is the scale or the criterion used?);

(b) To define a way this information can be exported from a given context to be used in another context (e.g. how can ratings or tags be exported and used by some recommender system?).

In the next paragraphs we will try to elaborate on these two issues.

```
<title>Attention stream of Jehad Najjar</title>
<feed>
  <title>Ariadne Indexation and Search Tool</title>
- <item>
    <title>Java lesson on inheritance</title>
    <type>html</type>
  - <events>
    - <event>
      - <action>
          <actionType>read</actionType>
        </action>
        <dateTime>2007-11-17T09:30:47.0Z</dateTime>
        <duration>120</duration>
      </event>
    - <event>
      - <action>
          <actionType>evaluate</actionType>
          <describingSchema>http://www.LRE.com/EvaluationValues</describingSchema>
        - <relatedData>
          - <entry>
              <name>Overall</name>
              <content>4</content>
            </entry>
          </relatedData>
        </action>
        <dateTime>2007-11-17T09:35:40.0Z</dateTime>
      </event>
    - <event>
      - <action>
          <actionType>tag</actionType>
        - <relatedData>
          - <entry>
              <content>Climate change</content>
            </entry>
          </relatedData>
        </action>
        <dateTime>2007-11-17T09:37:50.0Z</dateTime>
      </event>
    </events>
  </item>
</feed>
```

Fig. 2. XML example of CAM attention feed.

4 ANNOTATION SCHEMES

An important issue related to the representation and export of user feedback data sets from a given application environment is the declaration of the type(s) of feedback that is being collected and its format. For example, such an exported data set has to declare if the collected information is in the forms of ratings, reviews, or tags. In addition, it has to declare the exact structure and value spaces of the collected feedback. For instance, ratings may be collected upon one or more attributes (criteria), and using different rating scales. This can be particularly evident when examining evaluation or quality models for different application areas. In technology enhanced learning, we have found that rating of learning resources takes place using a variety of review schemes or instruments [5]. We call the definition of the structure and value

space of some user feedback type as an annotation scheme: this can refer to a particular model, set of criteria, or instrument that is used for collecting user feedback in a particular application environment.

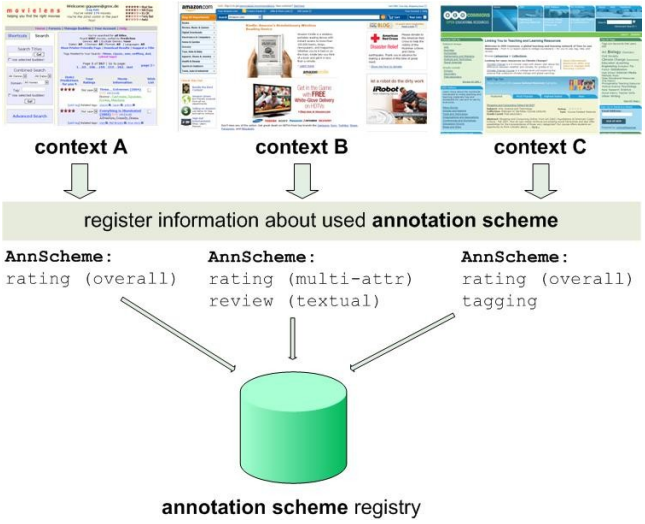

Fig. 3. Registering the annotation scheme used in different contexts.

In order to facilitate interoperability between different systems, and potential reusability of exported user feedback in other application environments, we argue that it is important to define the particular annotation scheme used. This is illustrated in the example presented in Figure 3. The movie recommender uses an annotation scheme that collects single-attribute ratings on movies and reflects users' overall satisfaction using a scale between 1 and 5. The e-commerce site, on the other hand, collects ratings on products. In this case, the ratings are collected using a scale from 1 to 7 upon different attributes (e.g. the quality of the product, its cost, and its availability). Additionally, textual product reviews by users are collected. Lastly, there is the educational portal that collects user feedback by asking users to rate their overall satisfaction with the item on a scale 1 to 7. Users can also provide their own tags to the items. Imagine that a recommender system would like to reuse the ratings on the same items (e.g. movies) that are collected by all three applications. In order to appropriately combine and transform existing ratings in a comparable format, the new system has to know (i) which annotation scheme has been used and (ii) its particular properties.

For the case under study, the solution of developing a registry of annotation schemes for learning resources has been chosen. In this way, it can be made possible for a learning resource collaborative filtering system [6] to collect ratings from different environments, and to refer to the annotation registry in order to appropriately combine or transform them. Similarly, existing textual reviews or tags can be collected and reused. For the Organic.Edunet project, a common annotation scheme registry will be developed, so that it is used as a reference point by all online access environments that collect user feedback. It is our intention to offer this registry openly for use by other initiatives in the domain of digital educational resources.

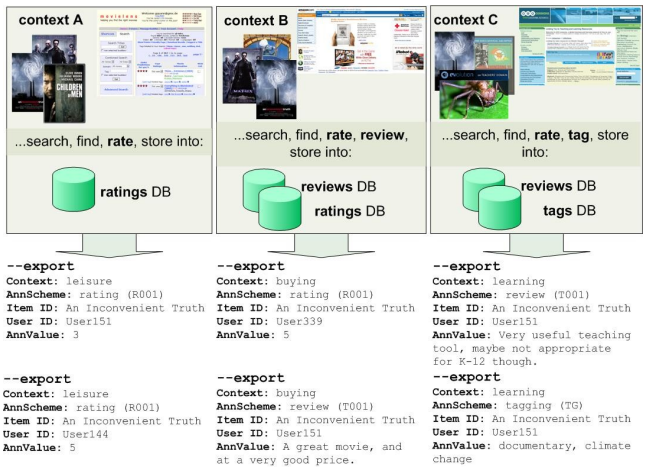

Fig. 4. Exporting annotations from different contexts.

5 COLLECTING DATA SETS

When the time comes for an application that collects user feedback to share this information with other applications, an appropriate transformation has to be followed. Figure 4 outlines how a data set could be transformed and exported in the conceptual example that we have used so far. The three different application environments that we consider are not simply dumping the user data in their databases with ratings, reviews or tags. They follow a more structured representation, and also export information about:

(a) the context in which this user feedback is collected, so that other applications can decide if and how they should reuse this particular feedback;

(b) the annotation scheme used, so that other applications can find out which annotation scheme has been used and appropriately transform the exported data into a format of its needs.

In our case study of annotating and recommending learning resources, this information is stored according to the adopted CAM schema (Section 3) into a central CAM repository. In this way, it is made available to all the different applications that can be developed to take advantage of this data.

For instance, Organic.Edunet will develop a collaborative filtering service for learning resources [7]. This service can search for user ratings that have been collected upon various resources from the CAM repositories of both MELT and Organic.Edunet, refer to the annotation scheme registry in order to decide how the collected ratings can be transformed into a comparable format, and use them as input for producing new recommendations. In a similar way, MELT can collect tags from MELT, Organic.Edunet and other educational portals (such as OERCommons, http://www.oercommons.org/), in order to use them in a cross-repository way (e.g. creating cross-repository tag clouds). Additionally, this type of data sets can be used to run evaluation experiments of new algorithms to be used in an educational context.

The requirement for the application environments that collect user feedback will only be to export their data sets in a CAM-compliant way, so that it is stored in a reusable manner in the centralized CAM repository. This can be either manually

performed at some points in the lifetime of the systems or through a regular harvesting procedure (e.g. using OAI-PMH, http://www.openarchives.org/).

6 PILOT IMPLEMENTATIONS

The ideas described in the previous section have been implemented in the initial version of a Social Navigation Module that aims to support recommendations of learning resources in federations of learning repositories. At the moment, the Module is being integrated with the collaborative filtering service of Organic.Edunet. It is also explored, how it can support recommendation in the context of the LRE portal (http://lreforschools.eun.org).

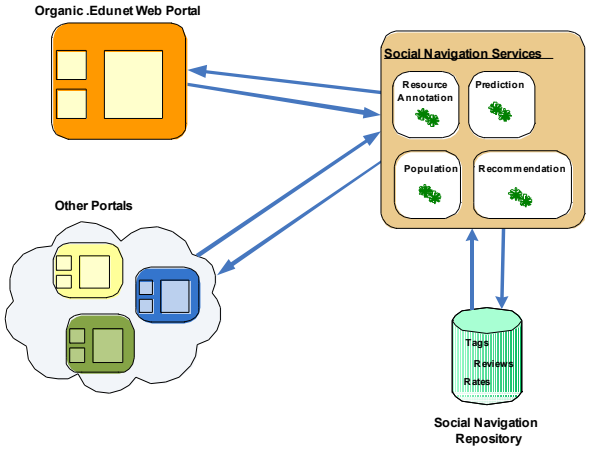

Fig. 5. Architecture of Social Navigation Module and how it supports Organic.Edunet.

Figure 5 illustrates how the Social Navigation Module is collecting user annotations from the Organic.Edunet Web Portal, and also provides recommendations based on these annotations. It also shows that in a similar way, other Web portals can be supported by invoking the relevant services. In addition, Figure 6 shows an example of how the different features of the Social Navigation Module can be provided through the interface of a Web portal. More specifically, it illustrates primitive interfaces for collecting tags, reviews and ratings from some user, and calculating a collaborative filtering prediction out of these annotations. Although this is a draft interface design that does not show how the features can be integrated into a fully operational system (e.g. Organic.Edunet Web Portal), it gives an indication of the way the Social Navigation Module can be invoked to support recommendation.

7 OTHER POTENTIAL USES

Apart from the implementation in the context of the two European projects that has been outlined above, the representation of user feedback in a structured, reusable and interoperable format can add value for other envisaged applications as well. Let us

consider the case of some user that is regularly using a number of different educational portals (e.g. OERCommons, MERLOT, http://www.merlot.org) in order to find, access, view or download learning resources that she finds listed there. Apart from using these resources for her own learning or teaching needs, the user can also provide feedback such as ratings, tags and pedagogical descriptions of usage.

Instead of storing her feedback in the server of the OERCommons, we could envisage the storage of this data in her own personal portfolio of resource annotations. In this way, feedback can be collected when she visits and annotates resources listed in other educational portals (such as MERLOT) or in other popular sites (like Amazon). When logging into the OERCommons portal again, her personal tag-cloud can be enriched with tags that she has provided in other environments that operate in an educational context. Or, when she logs into Amazon in order to get personal recommendations about appropriate textbooks for her teaching, the textual reviews that she has provided in the OERCommons portal and the ratings that she has provided on similar items in the MERLOT portal can also be considered.

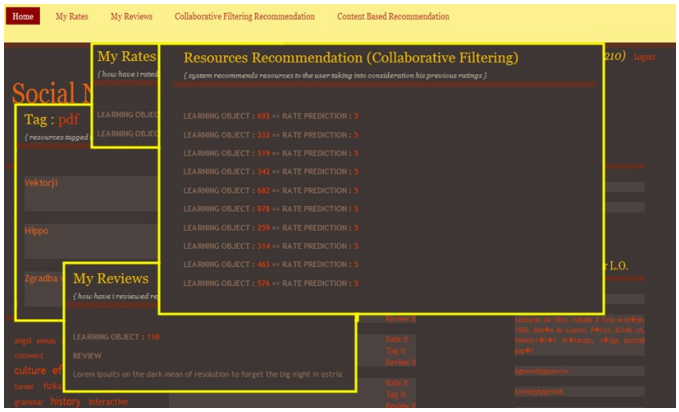

Fig. 6. Example interface of Social Navigation features.

8 CONCLUSIONS

Overall, we believe that the potential of collecting user feedback in a structured, reusable and interoperable (i.e. commonly agreed) format is high. In the context of MELT and Organic.Edunet, a Social Navigation Module is being developed based on the CAM-oriented approach presented in the paper. This Module aims to support recommendation of learning resources in federations of learning repositories." we currently implement the CAM solution examined above, and we expect to have some initial output for demonstration later in 2009.

However, several issues have to be addressed in order to apply such solutions in a larger extend. For instance, an important requirement will be how to uniquely identify the users when changing from one system to another (especially if they are using different user IDs) or the items when they are catalogued in different systems (especially when referring to items that do not have a specific URI). In addition, the adopted CAM solution is a rather user-oriented solution (since it captures attention streams of particular users) and collects very large information streams. The way that this will be implemented in scenarios such as the one examined in Section 7 has to be

explored, in order to achieve fast data transfer and on-the-fly calculation of recommendations.

The approach presented in this paper seems promising, but there is a number of issues that still have to be elaborated and explored. For example, more work has to be carried out on how the context in which user feedback is collected can be represented (especially as far as the "learning" context is concerned), and how this information is actually used. Furthermore, the pilot implementation of the approach in the Social Navigation Module can provide a test bench for an evaluation experiment that will verify the reusability of the annotations. Finally, the criteria upon which learning resources are assessed also affect the meaning of the recommendation. That is, it is different if learners assess their overall satisfaction from using a particular resource, in comparison to teachers assessing how this resource can be used as a teaching aid. This is another important reason for further exploring the dimensions used to represent the context and the preferences of users, when feedback is collected.

Acknowledgements. Part of the work presented in this paper has been funded with support by the European Commission (project No ECP-2006-EDU-410012 Organic.Edunet) and with a stipend from HS 100-vuotissäätiö.

REFERENCES

1. Herlocker, J., Konstan, J.A., Riedl, J.: An Empirical Analysis of Design Choices in Neighborhood-Based Collaborative Filtering Algorithms. Information Retrieval. 5 (2002)
2. Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating Collaborative Filtering Recommender Systems. ACM Transactions on Information Systems. 22(1) 5-53. (2004)
3. AttentionXML: AttentionXML specifications, Retrieved June 8, 2007. From http://developers.technorati.com/ wiki/attentionxml. (2004)
4. Najjar, J., Wolpers, M., & Duval, E.: Attention Metadata: Collection and Management. Paper presented at the World Wide Web 2006 Workshop Logging Traces of Web Activity: The Mechanics of Data Collection, May 23, 2006, Edinburgh, UK. (2006)
5. Vuorikari, R, Manouselis, N., Duval, E.: Using Metadata for Storing, Sharing, and Reusing Evaluations. In Social Recommendation: the Case of Learning Resources, in Go D.H. & Foo S. (Eds.) Social Information Retrieval Systems: Emerging Technologies and Applications for Searching the Web Effectively, Hershey, PA: Idea Group Publishing, 87-107. (2008)
6. Manouselis, N., Vuorikari, R., Van Assche, F.: Simulated Analysis of MAUT Collaborative Filtering for Learning Object Recommendation, in Proc. of the Workshop on Social Information Retrieval for Technology-Enhanced Learning (SIRTEL'07), 2nd European Conference on Technology Enhanced Learning, Crete, Greece (2007)