



Resources and users in the tagging process: approaches and case studies

Felice Ferrara¹, Ilaria Torre², Luigi Sarti³, Carlo Tasso¹, Antonina Dattolo¹, Stefania Bocconi³, Jeff Earp³

¹Università degli Studi di Udine, ²Università degli Studi di Genova, ³CNR

Keywords: social tagging, ontologies, user profiling, clustering, neighbour selection

Abstract

In this contribution we propose a comparison between two distinct approaches to the annotation of digital resources. The former, top-down, is rooted in the cathedral model and is based on an authoritative, centralized definition of the adopted mark-up language; the latter, bottom-up, refers to the bazaar model and is based on the contributions of a community of users. These two approaches are analyzed taking into account both their descriptive potential and the constraints they impose on the reasoning process of recommender systems, with special reference to user profiling. Three case studies are described, with reference to research projects that apply these approaches in the contexts of e-learning and knowledge management.

1 Introduction

The process of tagging resources is fundamental for building Web information systems, and it is a crucial task in environments designed to speed up the process of accessing knowledge, such as recommender systems, which are the subject of this work.

Previous works basically describe approaches of two different kinds: top-down and bottom-up. In top-down approaches, the tagging process is constrained by an *a priori* model for assigning metadata to resources. The model that guides the association of metadata may be based on well-known standards, such as LOM (Learning Object Metadata) and IMS-LIP (Instructional Management Systems - Learner Information Package), or alternatively on ontologies, which are machine readable structures aimed at modelling, organizing and representing a particular domain of knowledge (Gavrilova *et al.*, 2009).

It is important to stress that ontologies model relevant aspects of a particular domain from the perspective of a specific community of users. For this reason, factors such as the evolution of the community, the influx of new kinds of users, the introduction of new ways of interacting with the system all demand adequate flexibility, which is associated to the evolutionary power of the tagging model (Mika, 2005). A tagging model should be able to cope with these problems in order to support users effectively during the phase of information searching.

A tagging model should also offer multi-faceted representation of resources, something which can be very useful when dealing with a heterogeneous set of users belonging to distinct communities. The term “multi-faceted” aims to define a characterization of resources which can represent both denotative and, possibly, connotative meanings. Denotative and, mainly, connotative meanings can change according to the different viewpoints that different users may have of the same resource.

This requirement is better supported by a bottom-up strategy than by a top-down one. The most common and popular expression of bottom-up approaches is social tagging, which allows users to share and define personal classification by means of labels, known as tags.

Tagging mechanisms not only serve to classify digital resources, they also allow modelling of users’ information needs. This is particularly true when information is accessed via filtering and recommending systems, which need to model the knowledge, interests and goals of each user looking at the set of annotated resources.

The close connection between resource modelling and user profiling is borne out in the definition and characteristics of filtering and recommender systems, which are fundamentally based on matching user and resource models, defined

in several ways.

This work focuses on this matching functionality in top-down and bottom-up approaches applied in e-learning and knowledge management settings¹.

The paper is organized as follows: section 2 describes applications, characteristics and potentialities of top-down and bottom-up approaches with reference to the task of modelling users and resources; section 3 reports on three case studies which highlight how the tagging process can be integrated in traditional ontological mechanisms for modelling resources and also used to model user interests and characteristics for recommending resources; section 4 closes the paper by proposing some final considerations and cues for future work.

2 Resource classification and support for the processes of user profiling and resource filtering

Resources and users are the main components in filtering and recommendation mechanisms. In content-based recommender systems (Pazzani & Billsus, 2007), resources are usually annotated using a set of features that can be adopted for describing user interests as well. This way, each user can be related to the resources he/she is mostly interested in. However, it is also possible to relate users and resources by considering other features that are not based on content (namely features that do not explicitly express an interest), but that can also prove relevant for filtering and suggesting resources to the user. In the e-learning domain, examples of relevant features are learning goals, user knowledge and background, learning style, etc.

Conversely, in collaborative filtering approaches, the mechanism of recommendation is based on the identification of “similar” users. Resources that are relevant to some users within a cluster are suggested to other users belonging to the same cluster (Schafer *et al.*, 2007).

In the following we analyze how top-down and bottom-up approaches have been used to represent resources and users and also to support reasoning mechanisms capable of inferring user features that are then used in the recommendation process.

2.1 Classification and representation of resources

This section analyzes how the classification and representation of resources can be supported by means of top-down and bottom-up approaches.

¹ The e-learning and knowledge management fields are closely connected since they largely share principles, techniques and goals. Knowledge management techniques define structures to support the creation, access and sharing of knowledge, while e-learning focuses on learning processes both from an individual and a social point of view (Schmidt, 2005).

2.1.1 Top-down Approach

Various authors point out the advantages of using the cathedral model (Wong *et al.*, 2008), and especially an ontology, for representing and modelling knowledge. Ontologies feature unique definitions for describing concepts and expressions within a closed vocabulary. Classes, instances, attributes and axioms can all be adopted to model structural and hierarchical relations; they also permit automatic inference processes that can lead to the “discovery” of new knowledge (Gruber, 1993).

In this light, ontologies provide a suitable basis for developing a metadata model with elements derived from the concepts (classes) of the ontology itself.

Specifications such as LOM (IEEE 2002), for describing digital content, or IMS-LIP (2001), for user profiling, have proved useful for supporting interoperability between systems and for document exchange. While the XML versions of these standards are widespread, they do not have the expressive power of an ontology and do not allow flexibility in knowledge representation. The main difference between the two approaches is that while XML schemas are used for modelling XML documents, ontological languages are used for modelling knowledge. XML is a language for modelling data, while ontological languages model metadata (Nilsson, 2001). Subsequently, not all the knowledge expressed in an ontology can be easily replicated in a model based on the standards mentioned above.

For these reasons, ontologies are not just useful resources to draw on in the design phase. They can also be made available to the applications and services that the user interacts with run-time, so that inference mechanisms can be leveraged to enhance recommending and search functions. Some systems even give the user direct online access to the ontology via an interactive navigation interface (e.g. MACE, 2009)².

2.1.2 Bottom-up approach

The approach based on the bazaar model (Wong *et al.*, 2008) is in many ways complementary to the top-down approach; it offers a variety of perspectives on resources, and supports the evolution of the tagging process over time. This evolution is reflected in denotative and connotative changes regarding users, resources, vocabulary and practices, all of which is captured by the dynamic annotation process.

Typically, the collection of tags that a user assigns is defined as a personomy, and a set of personomies is called a folksonomy, namely a representation

² MACE (2009). Metadata for Architectural Contents in Europe. Last retrieved from <http://portal.mace-project.eu/> on Feb. 25, 2010.

that encompasses personal and global dimensions. User defined classifications, mainly aimed at satisfying personal needs, are shared globally: each user can browse the network of resources following tags applied by other users.

In e-learning and knowledge management, this model can be very useful for detecting not-institutional paths of research and study (serendipity) and for supporting those students who have not yet mastered the concepts, tools and methodologies of a particular field. By allowing users to annotate contents, different classification schemata can emerge, extending the number of paths that students can follow in order to find suitable resources.

Similar considerations have also been made in the context of accessing museum resources, where there is a need to address the terminological gap between official classifications of artworks and the language of visitors (Chan, 2007).

During the tagging process it is also possible to collect statistical information about the usage of tags for a specific user and the whole community; this is typical in social bookmarking tools. This way, folksonomies make it possible to create a sort of indexing of resources based on tag occurrences for each individual resource; typically, tags are graphically represented in a cloudlike cluster where font size conveys the relative popularity of a given tag, i.e. the bigger the font, the more popular the tag. Often, the most popular tags represent the denotative meaning of the resource (the one privileged by top-down classifications), while the bottom-up approach gives visibility to connotative and personal meanings.

Tags associated to resources can be used to support information search using traditional information retrieval mechanisms. However, the possibility of analyzing user tagging patterns opens up interesting prospects for personalized information access mechanisms, such as in recommender systems.

The tagging process has been used both in content-based recommender systems (Musto *et al.*, 2009; Mobasher *et al.*, 2008) and in collaborative filtering systems (Zanardi & Carpa, 2008; Nakamura *et al.*, 2007).

2.2 User profile and personalization

This section analyzes the potential that the two approaches offer for supporting user profiling and personalization mechanisms.

2.2.1 Top-down approach

In top-down approaches, the use of annotation as a mean to infer knowledge about users is exclusively based on user actions on the tagged resources. Conversely, in bottom-up approaches, it can be based on the user annotations

as well.

In this section we describe two ways that top-down annotations are frequently adopted for inferring knowledge about users and personalizing interaction.

The first way involves building an overlay model for mapping the domain description to user knowledge about that domain. This is one of the most frequent ways to profile users in the e-learning field (see for example Brusilovsky & Millan, 2007). An overlay model makes it possible to represent the competences a user already possesses (prerequisites) and needs to acquire (learning goals) by associating each knowledge level to each skill level. In the process of annotating learning resources, the skill concept is very useful for representing not just the knowledge domain but also the ability the user can develop by interacting with a given resource.

The second frequently adopted way is the possibility to propagate the user's interest in a resource (or knowledge about it) to the properties that describe that resource, and thus to all the resources annotated with those specific properties. Therefore, recommendation can be performed without the need for an explicit overlay model. For example, in CHIP (Aroyo *et al.*, 2007), resources are semantically annotated and users are required to rate their interest in a set of resources. Each resource has different properties, annotated with different tags. When the user rates a resource, this rating is propagated to its properties. This technique makes it possible to suggest resources that are annotated with such properties or equivalent ones.

2.2.2 Bottom-up approach

The top-down approach allows direct comparison between the behaviour of each user and the ontological description of the specific domain. By contrast, the bottom-up approach entails translation of the folksonomic structure, which lacks information about the relationships among tags, into a more well defined structure, for example by clustering the set of available tags (Schwarzkop *et al.*, 2007).

The process of structuring annotations not only allows to reuse techniques typical of the top-down approach, it also provides structural information regarding the set of users who annotate the resources. Thus individual user behaviour can be modelled in terms of both tag semantics and similarity with peers. In the former case, tags are associated to concepts defined in lexical ontologies, such as WordNet, or to concepts defined in other taxonomies, such as Wikipedia, in order to infer knowledge about the users' interests (Szomszor *et al.*, 2008). In the latter case, the target user profile is built by analyzing individual tagging behaviour and then comparing this to profiles of other peers in order to detect

users with similar interests (see the case study described in Section 3.2) or with similar socio-demographical characteristics (see the case study presented in the Section 3.3).

It is important to emphasize that while the top-down approach is more suitable for modelling the knowledge of a user during a learning process, the bottom-up approach is especially suitable for supporting an explorative phase and is very effective for improving knowledge sharing and peer tutoring.

3 Three use cases

The three use cases described in this section all have something in common: the use of top-down and bottom-up approaches to improve access to resources. Each follows a different strategy. The first case (section 3.1) shows how a top-down approach can support the sharing and classification of resources for e-learning, with some room for integrating the bottom-up approach. In the second (section 3.2), recommending mechanisms are built according to a bottom-up approach that features modelling of user interests. The third case (section 3.3) draws on users' socio-demographic data associated to annotated resources.

3.1 A top-down approach to describing educational resources

The first use case presents strategies that have been adopted in the context of Share.TEC³, a European project devoted to pre-service and in-service Teacher Education (TE). The goals of Share.TEC are to: build a system that aggregates metadata descriptions of TE-oriented digital resources produced Europe-wide; provide personalized, culturally-sensitive brokerage for retrieving relevant digital content; support the sharing of knowledge and practices within the TE community in Europe.

The semantic core at the heart of the Share.TEC system is a Teacher Education Ontology (TEO). The ontology covers concepts relevant to the domain of Teacher Education, with particular regard for aspects considered pertinent to the sharing of digital resources and practice among potential members of the Share.TEC community, namely teacher educators, teachers, academic/educational publishers and content developers.

The purpose of TEO within the Share.TEC system is to provide: pedagogical characterization of digital content; representation of user profiles and competencies; a basis for multilingual and multicultural functionality; support for personalized interaction with adaptive user applications; support for the implementation of recommending functions.

TEO comprises five distinct but interlinked branches:

³ SHAring Digital REsources in the Teaching Education Community, eContentplus programme (ECP 2007 EDU 427015) <http://www.share-tec.eu>

1. digital content - educational resources and artefacts closely related to the concept of “learning object” (Alvino *et al.*, 2007; Wiley, 2000);
2. actor and role - description of the system’s end-users in terms of professional profile, experience, personal preferences, etc. (Razmerita *et al.*, 2007; Paneva, 2006);
3. competencies – “statements that someone, and more generally some resource, can demonstrate the application of a generic skill to some knowledge, with a certain degree of performance” (Paquette, 2007);
4. context - for example the characteristics of the organizations in which Share.TEC’s end-users operate;
5. knowledge domain - taxonomic representation of knowledge areas related to Teacher Education, including pedagogical, technological and disciplinary fields (UNESCO - ISCED, 1997 - International Standard Classification of Education).

A top-down ontological approach (section 2.1.1) adopted for modelling general knowledge is integrated bottom-up with user-generated folksonomies and social tagging (section 2.1.2). These two processes are distinct but complementary: social tagging enriches the description of digital contents without directly affecting the ontology’s structure.

One of the key areas that Share.TEC addresses is the multicultural and multi-linguistic dimension associated with describing digital resources in a European context. When teacher educators search for digital resources, they often find themselves having to navigate a flood of results generated by generic search engines, and may have to cope with databases whose interface and metadata contain terms derived from an unfamiliar language and cultural model.

For this reason TEO adopts a two-level structure featuring a common layer and a series of specific ontologies, each capturing a particular cultural and linguistic context. This multi-layered ontological model provides a framework for the definition of Share.TEC’s metadata model. Similarly to the ontology, the metadata model comprises a common upper level (Common Metadata Model) and a series of language/context based derivations (Multicultural Metadata Model). The network of relations among these levels makes it possible to capture contextual, linguistic and semantic differences and to represent them in the Share.TEC system.

This approach offers users a number of advantages. Firstly, the system interface is available in the language of each consortium partner (Bulgarian, Dutch, English, Italian, Spanish, Swedish). Registered users are associated to ontology-based profiles whose main concepts and relations are linguistically and culturally contextualised and which are also represented in terms of competencies expressed within TEO (section 2.2.1). This means that users can search

for resources from a range of different linguistic and cultural settings and do so in their own language and using the terminology typifying their specific cultural model. In addition, inferencing capabilities have been developed that endow the system with added flexibility. This means, for example, that search results can be provided which, although not exactly matching a user's formal query specification, still satisfy their needs.

3.2 Bottom-up approach for interest based collaborative filtering

The second case study strictly focuses on the role of tagging for recommender systems. In particular, it describes a recommender system based on the idea that social bookmarks are a knowledge base that is built by users in order to satisfy some personal aims, and for this reason they represent a map of users' interests. Following this idea, it is possible to use the bottom-up approach to model both user interests (Section 2.2.2) and the relationship between a resource and an interest (Section 2.1.2).

Traditionally, social bookmarking applications provide users with tools for searching information that rank resources according to the popularity of the documents associated to the input tags. However, even though the popularity of a resource can be considered a good indicator of its quality, it does not take into account the heterogeneity of a network comprising students and teachers who each use tags to satisfy a personal and diversified set of information needs (both during classification and search phases).

Conversely, in collaborative filtering systems the task of finding a subset of people similar to the target user is fundamental for generating a list of recommendations.

In order to evaluate if two users are similar, collaborative filtering recommender systems have to compare feedback provided by the two users when they interact with the same resources or, at least, with related resources. For this reason, recommender systems achieve good results only when resources belong to the same information domain (Frankowski *et al.*, 2007). The traditional neighbour selection approach assesses the similarity between two users by counting the number of shared resources: the higher the number of shared resources, the higher the probability that the users will share other resources in the future. However, this approach fails when resources belong to different information domains.

This issue has been explored in (Baltrunas & Ricci, 2008), where the authors show that the accuracy of predictions significantly increases when the similarity among users depends only on resources connected to the current user interest. The quality of predictions generated by a recommender system can be enhanced by adopting a strategy capable of detecting both users' interests and

the relationships among resources and interests. The system can dynamically select the neighbourhood by comparing users only by drawing on the subset of those resources that are connected to the current user interest.

Following this idea, an interesting line of research is to adopt user defined classifications for mapping users' interests and generating a list of personalized recommendations.

To tackle the absence of semantic connections we need to infer relationships among tags. Specifically, we build a matrix where rows and columns are respectively associated to tags and resources; the cell (i,j) counts the number of times that the i -th tag has been associated to the j -th resource. The cosine similarity is then used to quantify the similarity among two tags (two rows of the matrix). These similarity relationships are used to group users' tags into clusters: given a user, one of his/her interests is defined by a cluster of tags and the set of resources he/she associated to tags in the cluster.

Given a user interest, the list of recommendations can be generated by taking in account only people who bookmarked resources using labels connected to tags used by the target user for referring to the specific interest. In other words, the neighbour selection phase is defined by considering only tags and resources connected to the current user interest.

In order to detect all tags connected to the current user interest, the set of tags applied by the target user is further expanded to include other similar tags.

The expanded set of tags can be used for:

- filtering the set of potential neighbours. Only users who applied tags related to the current user interest are considered;
- filtering the set of potential recommendable resources. Only resources which have been associated to tags related to the current user interest are considered.

The traditional collaborative filtering approach can be used on this filtered set of data: a score is assigned to each potential neighbour to quantify his/her similarity to the target user according to the number of shared resources; a score is assigned to each potential recommendable resource, according to the similarity of neighbours who bookmarked it, in order to quantify the relevance of resources. In the end, the most relevant resources are suggested to the target user. More detail about the recommender system can be found in (Dattolo *et al.*, 2009).

3.3 Bottom-up approach for peer-tutoring

The last case study regards a recommender system that, similarly to the pre-

vious one, uses social bookmarks as a repository of resources and a mechanism of neighbour selection to identify similar users. However, the user dimension considered for personalizing the results are different in the two case studies, as well as the modalities adopted for acquiring them.

Considering the classification in section 2, this case study outlines the potential of social tagging, applied to social bookmarking, for selecting and indexing resources according to the users' point of view. Users are considered as complex and multidimensional entities. Each user dimension, or characteristic, influences the annotations users produce. Thus, correspondingly, it can be useful to take these characteristics into account when filtering and recommending resources. The system described in this case study focuses on this aspect, since personalization of results is obtained by matching a set of features characterising the query-performing user with those characterising users who annotated the resources.

The basic idea is that, if a user is looking for information on a given subject like "how to configure a router", for instance, he/she will prefer to find resources bookmarked and tagged by people who share the same problem and who have the same background, so that the explanations fit her ability to understand them. In satisfying this need, the bottom-up approach clearly offers support both for the classification of resources and for the identification of similar users who bookmarked, as relevant, a resource on the subject of the query.

This approach is especially relevant in the e-learning domain, and is supported by studies about the effectiveness of peer tutoring (Westera, 2007). These studies highlight that suggestions from fellow-students about materials for learning tasks are usually taken into greater consideration than suggestions from teachers.

The implemented prototype (see details in Torre, 2009) uses Delicious, and its social bookmarks, as a repository of tagged resources. Given a query made of one or more tags, the system returns the set of resources (bookmarks) that satisfy the query, and ranks the results by considering the tag occurrence. To this end, it uses the TF.IDF technique typically used in information retrieval. This technique is applied to the retrieval of social bookmarks as follows. For each resource whose tags match the query tag(s), the TF (Term Frequency) is calculated as the ratio between the number of occurrences of the query tag and the total number of tags applied by users to that resource. The IDF (Inverse Document Frequency) determines the relevance of the tag to the resource compared to its relevance to all the other resources. It is computed as the ratio between the total number of resources and the number of resources annotated with the query tag and taking the logarithm of the quotient. Finally, TF and IDF are multiplied and the value obtained represents a weight for the resource (bookmark), given the query tag.

Bookmarks ranked according to the weights obtained with the TF.IDF technique are then re-ranked in a personalized way by considering the match between the characteristics of the querying user (e.g. age, gender, etc.) and the characteristics of the users who tagged the bookmarks. This phase is made possible by the bottom-up approach, where each bookmark is associated to the users who bookmarked that resource and annotated it with a set of tags.

The users who bookmarked the resources are characterized by a set of attributes that, altogether, form the user identity: name, surname, age, gender, interests, job, school, etc. When the user makes his/her query, he/she can specify which attributes the system is to take into account for personalizing the results. For example, a student who is looking for resources about the PHP language to complete an exercise could make a query in order to retrieve, for example, a reference to a PHP library, specifying a preference for resources that have been bookmarked by students who share the same course of study.

The critical issue is how to acquire knowledge about the user who inserted and tagged the bookmarks. Delicious displays very little data about registered users. Thus, the problem is how and where to acquire further user data.

The approach followed to tackle this problem has been to develop a discovery service that works as a crawler of identities in social networks (cs-ids.di.unito.it). If a user is identified in one or more social networks, with a certainty factor over a given threshold (0.7 in the experimental evaluation) the profiles on the various social networks are combined and the new data discovered can be used to infer further knowledge about the user. Details about the crawler can be found in (Carmagnola *et al.*, 2009).

Exploiting this service, the recommender system can thus be used to build a sort of virtual class made up of virtual “fellow-students” selected from the whole set of Delicious users.

The current repository uses just one bookmarking system (Delicious) as a repository, but the system is conceived to query multiple systems in parallel so as to broaden the repository used for the query.

Conclusions

This work has analyzed the capabilities, possibilities and perspectives of tagging for supporting resource classification and user profiling, the two key tasks for developing recommender systems.

In particular, section 3.1 presented a study based on the top-down approach for supporting the classification of learning objects. The top-down approach offers marked descriptive power but incurs high costs in terms of knowledge building and maintenance. Conversely, the bottom-up approach distributes the effort of the classification task among users, who can label resources without

observing strict constraints; by the same token, however, it leads to difficulty in effective automatic analysis of the classification terms adopted.

The bottom-up approach opens up interesting opportunities for supporting the classification of users and resources. Sections 3.2 and 3.3 showed how the bottom-up approach can be used for recommending resources by looking at similar users (neighbour selection) and by considering respectively the similarity over both interests, inferred from tags applied by users, and socio-demographical data, imported from social Web environments.

Future analysis will focus on evaluation and comparison of the proposed methodologies using feedback from students and teachers in order to assess capabilities and possible integrations between top-down and bottom-up models.

Top-down and bottom-up approaches are, in several aspects, complementary, as shown by studies conducted by Macgregor and McCulloch (2006). Chan's studies showed that an ideal classification model should combine official and social classification (Chan, 2007). Hybrid models that seek to combine top-down and bottom-up approaches have been analysed in previous works (Van der Sluijs & Houben, 2008) and the case study proposed in section 3.1 points in this direction. Nevertheless further evaluation is needed to assess the descriptive power of the model.

REFERENCES

- Alvino S., Forcheri P., Ierardi M.G., Sarti L. (2007), *Un modello per la connotazione pedagogica di LO*, Rapporto tecnico IMATI-CNR n. 13/2007.
- Baltrunas L., Ricci F. (2008), *Locally adaptive neighborhood selection for collaborative filtering recommendations*, in: AH '08: Proceedings of the 5th International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, 22–31, Berlin, Heidelberg, Springer-Verlag.
- Brusilovsky P., Millan E. (2007), *User models for adaptive hypermedia and adaptive educational systems*, in: Brusilovsky P., Kobsa A., Neidl W.(eds.): *The Adaptive Web*, 3-53, Berlin Heidelberg New York: Springer-Verlag.
- Carmagnola F., Osborne F., Torre I. (2009), *Cross-Systems Identification of Users in the Social Web*, in Proc. of the IADIS International Conference WWW/Internet, 129-134, Rome, Italy.
- Chan S. (2007), *Tagging and searching? serendipity and museum collection databases*. In *Museums and the Web*, Archives Museum Informatics.
- Dattolo A., Ferrara F., Tasso C. (2009), *Neighbor selection and recommendations in social bookmarking tools*, in Proc. of the Workshop on AI & E-Learning, within the Conference of the Italian Association for Artificial Intelligence (AI*IA), 27-36, Reggio Emilia.

- Gavrilova T., Puuronen S., Alisova M., Petrashen, E. (2009), *First steps in Ontology Development: Knowledge Portal for Software Testers*, in: Dicheva D., Nikolov R., Stefanova, E. (eds.), *Software, Services & Semantic Technologies*, First International Conference, 192-199, Sofia, Bulgaria.
- Gruber T. (1993), *A translation approach to portable ontology specifications*, Knowledge Acquisition.
- IEEE Learning Technology Standards Committee (2002), *Learning Object Metadata*, Final Draft Standard. IEEE 1484.12.1-2002.
- IMS LIP (2001), *IMS Learner Information Package*, <http://www.imsproject.org/aboutims.html> (verificato il 28 Febbraio 2010).
- Macgregor G., McCulloch E. (2006), *Collaborative Tagging as a Knowledge Organisation and Resource Discovery Tool*, Library Review.
- Mika, P. (2005), *Ontologies are us: A unified model of social networks and semantics*, in: Gil Y., Motta E., Benjamins R.V., Musen M. (eds.), *Proceedings of the Fourth Int. Semantic Web Conference (ISWC)*, 122-136, Galway, Ireland.
- Musto C., Narducci F., De Gemmis M., Lops P., Semeraro G. (2009), *Star: a social tag recommender system*, in: *Proceedings of the ECML/PKDD 2009 Discovery Challenge Workshop*.
- Nakamoto, R., Nakajima, S., Miyazaki, J., Uemura, S. (2007), *Tag based contextual collaborative filtering*, *Journal of Intelligent Information Systems* 34(2).
- Nilsson, M. (2001), *The semantic web: how RDF will change learning technology standards*. <http://zope.cetis.ac.uk/content/20010927172953>, (verificato il 28 Febbraio 2010).
- Paneva D. (2006), *Ontology-based Student Modeling*, in: *Proceedings of the Fourth CHIRON Open Workshop Ubiquitous Learning Challenges: Design, Experiments and Context Aware Ubiquitous Learning*, 17-21, Turin, Italy.
- Paquette G. (2007), *An Ontology and a Software Framework for Competency Modeling and Management*, *Educational Technology & Society*, 10 (3), 1-21.
- Pazzani M., Billsus D. (2007), *Content-Based Recommendation Systems*, in: Brusilovsky P., Kobsa A., Nejdl W. (eds) *The Adaptive Web*, 325–341, Springer-Verlag, Berlin, Germany.
- Razmerita L., Angehrn A., Nabeth, T. (2003), *On the role of Actor models and Actor modelling in Knowledge Management Systems*, in: *Proceedings of HCI International Conference*, Greece, 2003.
- Schafer, J. B., Frankowski, D., Herlocker, J., & Sen, S. (2007), *Collaborative Filtering Recommender Systems*, in: Brusilovsky P., Kobsa A., Nejdl W. (eds), *The Adaptive Web*, 291-324, Springer-Verlag, Berlin, Germany.
- Schmidt A. (2005), *Bridging the Gap Between Knowledge Management and E-Learning with Context-Aware Corporate Learning*, *Professional Knowledge Management*, 203-213.
- Schwarzkopf E., Heckmann D., Dengler D., Kröner A. (2007), *Mining the structure of tag spaces for user modelling*, in: *Proceedings of the workshop on data mining for user modeling*, at the International Conference on User Modeling (UM '07),

- 63–75, Corfu, Greece.
- Shepitsen A., Gemmell J., Mobasher B., Burke R (2008), *Personalized recommendation in collaborative tagging systems using hierarchical clustering*, in: Proceeding of the 2nd ACM International Conference on Recommender Systems, 259-266, Lausanne.
- Torre I. (2009), *Social Tagging in eLearning. Recommendations from Virtual Fellow-Students*, in Proc. of the Workshop on AI & E-Learning, within the Conference of the Italian Association for Artificial Intelligence (AI*IA), 27-36, Reggio Emilia.
- van der Sluijs K., Houben, G. (2008), *Relating user tags to ontological information*. In: Proce. of the Workshop on Theories and Applications of Ubiquitous User Modeling, at the Int. Conf. of Intelligent User Interfaces (IUI '08), Canary Islands, Spain. <http://www.uni-koblenz.de/confsec/IUI/Proc./data/workshops.html>.
- Westera W. (2007), *Peer-allocated instant response (pair): Computational allocation of peer tutors in learning communities*. Journal of Artificial Societies and Social Simulation, 10 (2).
- Wiley D.A. (2000), *Connecting learning objects to instructional design theory: A definition, a metaphor, and a taxonomy*, in Wiley D.A. (ed.) The Instructional Use of Learning Objects, Association for Instructional Technology.
- Wong D., Shephard K.L., Phillips, P (2008), *The cathedral and the bazaar of e-repository development: encouraging community engagement with moving pictures and sound*, ALT-J Research in Learning Technology, 16 (1), 31-40.
- Zanardi V., Capra L. (2008), *Social ranking: Finding relevant content in web 2.0*, in: Proceeding of the 2nd ACM International Conference on Recommender Systems, 51-58, Lausanne.