

Tag-based User Modeling for Social Multi-Device Adaptive Guides

Francesca Carmagnola, Federica Cena, Luca Console, Omar Cortassa,
Cristina Gena, Anna Goy, Ilaria Torre
Dipartimento di Informatica, Università di Torino
Corso Svizzera 185; 10149 Torino, Italy
E-mail: {carmagnola, cena, console, cortassa, gena, goy, torre}@di.unito.it

Andrea Toso, Fabiana Venero
CSP Innovazione nelle ICT s.c.a.r.l., Torino, Italy
E-mail: {toso, venero}@csp.it

Abstract.

This paper aims to demonstrate that the principles of adaptation and user modeling, especially social annotation, can be integrated fruitfully with those of the Web 2.0 paradigm and thereby enhance in the domain of cultural heritage. We propose a framework for improving recommender systems through exploiting the users tagging activity. We maintain that Web 2.0's participative features can be exploited by adaptive web-based systems in order to enrich and extend the user model, improve social navigation and enrich information from a bottom-up perspective. Thus our approach stresses social annotation as a new and powerful kind of feedback and as a way to infer knowledge about users. The prototype implementation of our framework in the domain of cultural heritage is named iCITY. It is serving to demonstrate the validity of our approach and to highlight the benefits of this approach specifically for cultural heritage. iCITY is an adaptive, social, multi-device recommender guide that provides information about the cultural resources and events promoting the cultural heritage in the city of Torino. Our paper first describes this system and then discusses the results of a set of evaluations that were carried out at different stages of the systems development and aimed at validating the framework and implementation of this specific prototype. In particular, we carried out a heuristic evaluation and two sets of usability tests, aimed at checking the usability of the user interface, specifically of the adaptive behavior of the system. Moreover, we conducted evaluations aimed at investigating the role of tags in the definition of the user model and the impact of tags on the accuracy of recommendations. Our results are encouraging.

Keywords: Web 2.0, tag, annotation, user model, recommendations, cultural events

1. Introduction

The aim of this paper is to show that the combination of the principles of adaptation and Web 2.0 opens new opportunities for services related to cultural heritage. In the paper, in particular, we will present a framework for such an integration using cultural heritage as a reference domain both in the design and in the application of the framework.

© 2008 *Kluwer Academic Publishers. Printed in the Netherlands.*

Web 2.0 (OReilly, 2005) is a term commonly used to identify a second generation of Internet-based services that allow people to collaborate and share information online. It aims at decentralizing the creation/editing/organization of contents. In fact any user can become an author of contents, can make these contents available to all users (or groups of users) and can create (and share) her own personal view of the content defining personal criteria to aggregate and link them (Marlow et al., 2006). Users can annotate contents by adding keywords (tags) or additional information. This is the basis for creating folksonomies, bottom-up aggregations of contents that share a same tag (folksonomies work in an opposite way from taxonomies, top-down set-ups). Tags and folksonomies can be personal or shared. In this way, annotation and user participation also lead to the creation of communities and networks of users with similar or related interests or similar ways of approaching the content of a site.

These principles are particularly relevant for services related to *cultural heritage*. The notion of cultural heritage links two concepts: heritage and culture. The term “heritage” refers to local landscape, architecture and artefacts that make a particular place unique. The term “cultural” puts less emphasis on the place and, more generally, regards knowledge, patrimony, and local traditions of a territory. Therefore, while the heritage is something which can be seen as static, the notion of culture is more dynamic since it is bound to people living in specific places. In this sense, we term “cultural heritage resources” both the physical heritage and the intangible activities related to it, like events promoting or/and taking place in heritage places, stories, anecdotes and legends about the places, traditions linked to the place, etc.

Promoting (and living) the cultural heritage of a territory means giving information to people (and sharing information) about the heritage of a territory and giving voice to the events celebrating the culture of such a territory. Nowadays such a promotion can exploit the communicative power of the Web. Most areas - so called “territories” - have a web site that promotes the knowledge of their cultural heritage. These sites are typically managed by official agencies or public organizations. The resources are classified by specialists (Trant, 2006), on the basis of formal classification systems (Chan, 2007). Thus, the user has a passive role and she can merely access the content provided by others in a top-down perspective.

This perspective seems to be quite limited due to the dynamic nature of cultural heritage. Thus, such a perspective could benefit from the support of people who live or visit the territory, who are the effective recipients of cultural heritage resources. In this direction, the involvement of users in the promotion of cultural heritage resources finds a support

in the Web 2.0 vision and in its emerging techniques. In particular, Web 2.0 can enhance functionality of a web-based system, enriching the traditional web approach with other features such as (Sen et al., 2006):

- *Participation in the process of content creation.* Users can *i)* add new resources; *ii)* add new information about a resource; *iii)* add further detail about existing items; *iv)* provide comments and judgments to a resource. This leads to an increase of the *quantity* of information related to the cultural heritage (for example, they can insert events or information about an event celebrating the cultural heritage, such as a description or logistic information), and/or the *quality* of information, making existing information more precise (e.g. correcting misleading data). This can also lead to increasing the “narrative dimension” of the resources (Gorla et al., 2003), for example by enriching them with stories and anecdotes related to their history. Providing judgments can be an incentive for stimulating other users in experiencing a cultural heritage resource. For example, after attending an event, a user can express her satisfaction or disappointment about it. Each user can thus provide a personal point of view of cultural heritage contents with respect to information provided by experts. Having the user’s point of view in the description of an item can help other users in finding the right information (Chun and Wyman, 2006).
- *Self-organization of content.* Users can store and collect information for later reference. In addition to traditional client-side bookmarking, Web 2.0 enables users to add tags to the bookmarked information and to use such tags to retrieve it later from any connected computer. Tags, moreover, support the definition of folksonomies, which are bottom-up personal aggregations of resources offering a personalized way of navigation.
- *Socialization and collaboration.* Users aim at sharing contents with friends. Social annotation systems enable users to share their annotations and tags to a particular cultural resource (Marlow et al., 2006). A user can also exploit tags inserted by the community of users and the resulting social folksonomies to navigate content differently from the methods of official categorization.

Web 2.0 offers users a support to actively experience the cultural heritage of the place they live in or they are visiting. It also provides means to involve users in the promotion of such cultural heritage. Web 2.0 can also encourage the proliferation of social communities that

share interests about specific cultural heritage subjects (for example, a community of users interested in overhauling a historical place). In this way, Web 2.0 can provide support to keep the cultural heritage alive, stimulating people to live it actively.

However, Web 2.0 alone does not provide mechanisms powerful enough to cope with the problem of managing huge amounts of contents. Potentially, there are contents at different levels of detail (e.g. from high-level descriptions to very detailed ones suitable for expert users) and contents from different viewpoints (e.g., description of a resource focusing on historic or artistic view). *Adaptive web systems* (Brusilovsky and Maybury, 2002), (Brusilovsky et al., 2007) have the goal of tailoring the selection of resources, the presentation of contents and the navigation structure to the model of each specific user and context of usage. This can be particularly useful within a Web 2.0 application.

We claim that the combination of the principles of Web 2.0 and adaptation based on user models can be very fruitful. Adaptation can receive a new impulse from a Web 2.0 orientation and, in turn, a Web 2.0 service can benefit from adaptation. Here are some benefits: *i*) the active participation of a user, typical of Web 2.0, can provide a lot of information that can be used by an adaptive system to define more precise user models; *ii*) user models can be used to create communities of users with common interests; *iii*) user models can be used to provide a suitable and personal support for the annotation of contents, (e.g. to suggest tags); *iv*) adaptive techniques can help users in navigating and finding the preferred information in the huge amount of data inserted by a community of users; *v*) adaptation can support the use of multiple devices and multiple contexts of interaction of a Web 2.0 application.

We analyze this beneficial combination in an accurate way, focusing on adaptive Web 2.0 services related to cultural heritage, which are a perfect case study for investigating the integration of Web 2.0 and adaptation and to highlight and test its benefits. In particular, we propose a framework, called PAPERRE (PARTicipative PERSonalized REcommender), for the integration of adaptive web-based systems and Web 2.0. We first present an architecture for the integration (Section 2) focusing on Web 2.0 enhanced user modeling (Section 3). Afterwards we present an instantiation of PAPERRE introducing iCITY (Section 4), an application providing information about cultural resources and events promoting cultural heritage in the city of Torino. Finally, Section 5 presents some evaluation results, Section 6 discusses the relevant related work, and Section 7 provides some discussions and conclusions.

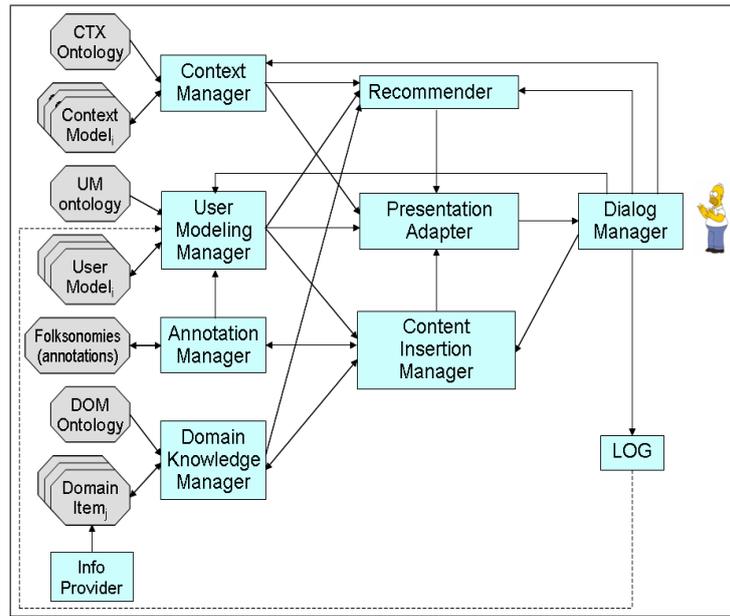


Figure 1. Abstract model of the architecture of the PAPER framework.

2. Architecture of the PAPER framework

In this section we present an architecture for supporting the integration of Web 2.0 and adaptation, discussed in the previous section. This architecture, depicted in Figure 1, represents a conceptual model, which describes the main modules cooperating within the PAPER framework to provide the requested functionality, together with the communication channels among such modules.

It is important to stress that this architecture represents an abstract point of view since it refers to the framework level. This means that it abstracts from any detail concerning both the actual possible implementation and the instantiations on specific domains. On the basis of this conceptual (abstract) view, in fact, each module can be instantiated by different concrete components, provided that they implement (in some way) the functionality defined for that module.

Another important characteristic of the architecture presented here is that the knowledge bases are completely declarative and they are accessed by dedicated modules. Thus, when a module needs to access a knowledge item, it simply interacts with the corresponding dedicated module, which is in charge of handling that knowledge base, without being aware of the details concerning knowledge representation.

Some of the activities performed by the modules can take place during the run-time of a system that instantiates the framework, while

some others are designed to be performed only off-line. In the following, when describing such activities, we will specify if they take place at runtime or off-line, and we will explain our choice. Before presenting each module, we will provide a description of the structure of the knowledge bases. Obviously, the content of the knowledge bases is domain dependent, and thus at the framework level, we can only outline their roles. We will provide some examples of a possible instantiation of these knowledge bases in Section 4.5.1, when an instantiation of the PAPERRE framework on the domain of cultural events (the iCITY prototype) is presented.

Domain (DOM) Ontology. This ontology defines the concepts needed to describe the items in the domain. At the framework level, we do not make any assumption about the content and the format of this ontology.

Domain Items and Info Provider. The Domain Items knowledge base contains the instances of the DOM Ontology. It can be viewed as a set of item descriptions based on the language defined by that ontology. This knowledge base is fed in by an Info Provider that can be implemented in different ways. Namely, various agents can extract information from an up-to-date database a human operator, a software module, or a module that automatically extracts information from RSS feeds, as in our prototype implementation iCITY (Section 4.5).

Context (CTX) Ontology. This ontology defines the properties needed to describe the context of use. The concept of context, analyzed in detail by (Abowd and Mynatt, 2000; Dey and Abowd, 2000), is very broad and it is hard to come up with one single definition. Almost everything can be considered as context of use. For example, there is the physical context that includes user location, information about the environment conditions (light, noise, temperature, time of the connection, walking/driving speed, ...), etc. There is the social context that includes features like the social community or group the user belongs to, the task she is performing and the relation with people close to her while she is interacting with the service, etc. Many researchers include also information about the user within the notion of context. In the PAPERRE framework, we decided to leave the user model as a separate type of information.

In any case, the choice about which context features should be modeled is application-dependent, and is made by the system designer when she instantiates the framework, in order to build an application based on it.

Context Model. The Context Model is a knowledge structure that represents the current context of use - i.e. the situation of the current

user interacting with the system (e.g. her position, the device she is using, the light and noise conditions, and so on). The Context Model provides specific values for the properties defined in the CTX Ontology.

User Modeling (UM) Ontology. This ontology defines the concepts needed to describe user profiles. The properties that can be included in user profiles can be features derived from registration data (e.g., age, job, etc.), features representing an *overlay model* (Brusilovsky, 1996) of the domain (DOM Ontology), inferred features (see Section 3). An overlay model provides for each domain feature an estimate of how much the user is familiar with, or interested in, that feature. Hence the overlay model can be seen as a set of pairs $\langle domainFeature, value \rangle$, where *value* can be discrete or probabilistic. As described in Section 3, in order to take advantage of the enhanced user modeling based on the user's tagging activity, the user model should include the dimensions related to such an activity. These dimensions are described in Section 3.

User Models. A User Model is a knowledge structure that represents the profile of a single (registered) user. The dimensions in which the user description is structured refer to concepts defined in the UM Ontology.

Folksonomies. *Annotations* include tags and comments and are the user contribution to the system knowledge. In this sense, they represent the users' view of the system contents - i.e. the *folksonomies*. Notice that here we use the term *folksonomies* in a broader sense with respect to its standard meaning, since we include in it, not only tags, but all information coming from the user, including comments. *Tags* can be freely inserted by the users, or they can be selected from a vocabulary associated to the DOM Ontology (see, for instance, the tag management in the iCITY prototype, described in Section 4.5.1). Similarly, also *comments* can be free texts, or they can be "controlled" - i.e. selected within a predefined (eventually structured) set. Annotations represent links between domain knowledge and knowledge about the user, since each annotation refers to a specific domain item and has an author (user) (Chun and Wyman, 2006).

Below we outline the functionality of each module as well as the interactions required in order to provide such a functionality.

Context Manager. This module represents the interface between the system and the knowledge about the context of use. When other modules (e.g., the Recommender, or the Presentation Adapter) need information about the context, they invoke the Context Manager (see the arrows from the Context Manager to the Recommender and the Presentation Adapter in Figure 1). The Context Manager manages two

knowledge bases. First, it manages the CTX Ontology, which represents a “static” knowledge base, defined at design-time (i.e. when the framework is instantiated on an application domain). For this reason the interaction between Context Manager and CTX Ontology is read-only (one-way arrow in Figure 1). Second, it manages the model of the current context of use, accessed by the system (implementing the PAPERRE framework) at run-time, in order to provide adaptation.

In iCITY (see Section 4.5), the Context Manager has only this simple role. However, from the abstract point of view of the PAPERRE framework, it would be possible to equip the Context Manager with reasoning capabilities in order to infer more complex context features. For this reason, the interaction between Context Manager and Context Model can take place in both directions (read-write; see Figure 1) - i.e. the Context Manager can also update the current Context Model.

Domain Knowledge Manager. The Domain Knowledge Manager provides other modules with the knowledge about the domain. It manages the DOM Ontology and the set of domain items. Similarly, the DOM Ontology is defined at design-time and is read-only, while the domain items knowledge base can be updated, not only by the Info Provider, but also by the user (through the Content Insertion Manager; see below).

Annotation Manager. The Annotation Manager provides other modules with an access to the annotations by handling the knowledge bases storing tags and comments. As already mentioned, annotations represent folksonomies - i.e. the users’ view of the system contents. Differently from ontologies, folksonomies are not defined at design-time by the system designer, but they are incrementally built by users. For this reasons, the interactions between the Annotation Manager and the Folksonomies (Annotations) knowledge base is read-write (see double-arrow in Figure 1). Moreover, the Annotation Manager provides the User Modeling Manager with information about user tags (single arrow in Figure 1).

User Modeling Manager. This module represents the interface between the system and the knowledge about the users, and in this sense it is analogous to the other dedicated modules. Every time a module needs to get any information about the user, it asks the User Modeling Manager. It manages two knowledge bases: the UM Ontology (defined at design time and thus read-only) and the current User Model. Besides this basic functionality, the User Model Manager can provide a more complex behavior that can be supported by sub-modules performing different types of reasoning:

- The *Initialization sub-module* initializes a new User Model by means of inferences based on the user registration data. This sub-module makes use of data provided by the user in the registration form in order to infer user features such as interests, preferences, attitudes, and so on. We should notice that both the kind of data elicited in the registration form and the dimensions in the User Model are application dependent and are defined at design-time.
- The *Behavior-Tracking sub-module* reasons about the data concerning the interaction of the user with the system in order to learn user’s interests and preferences that are used to update the User Model. These data items include “standard” user actions (like link selections, bookmarking, etc) as well as information about the user annotation activity (i.e. insertion of tags or comments), provided by a log file and by the Annotation Manager. The reasoning performed by this sub-module takes place off-line– i.e. the User Model is updated only when the user logs out, and the changes are available in next sessions. The reason for this choice is to avoid changes in the system behavior based on changes in the current User Model, which could puzzle the user; see (Nielsen, 1999).
- The *Similarity sub-module* calculates the similarity between user models on the basis of some similarity measure.

An example of a specific implementation of the different reasoning activities performed by the User Modeling Manager sub-modules will be presented in Section 4.5.2.

Recommender. This is a typical “reasoning” module. From the abstract point of view we are discussing here, this reasoning activity performed by the recommender should take input including information about the domain, the user and the current context of use. This reasoning activity should provide output including an evaluation of the domain items; this evaluation represents the personalized suggestions of the system. The input information is supplied by the Domain Knowledge Manager, the User Modeling Manager, and the Context Manager (see arrows from these modules to the Recommender in Figure 1). The output is an XML object that is provided as input to the Presentation Adapter.

The conceptual architecture does not commit on the specific techniques that could be used to implement the reasoning capabilities of the Recommender. However, from an abstract point of view, the reasoning activity is based on associations that match user characteristics and context features with properties of the domain items. Moreover, it

produces, as a result, a ranked list of domain items tailored to the user preferences and context of use. This list represents the *personalized content* produced by a system based on the PAPERRE framework.

Presentation Adapter. The Presentation Adapter is another typical “reasoning” module. While the Recommender adapts the content to the user and context features, the Presentation Adapter personalizes the presentation of such a content. It takes as input information about the user and the current context (provided by the User Modeling Manager and the Context Manager) together with the output of the Recommender. It provides a presentation by deciding which pieces of information should be included (longer or shorter texts, more or less detailed descriptions, and so on), selecting suitable linguistic forms and multimedia objects to be included. Moreover, the Presentation Adapter selects the proper layout on the basis of the user preferences and context characteristics. This last function is particularly important in order to tailor the user interface to the device used to connect to the system. The output of the reasoning activity of the Presentation Adapter is an enriched XML object, coupled with an XSL stylesheet.

Finally, the Presentation Adapter generates the user interface by applying the selected XSL stylesheet to the XML object representing the personalized content/presentation. The output of this transformation can be any XML-based document. For instance, in our prototype iCITY (see Section 4.5) the transformation generates an XHTML document (linked to a CSS stylesheet) for desktop PCs or laptops, and an XHTML MP¹ (linked to a CSS stylesheet) for mobile devices.

Content Insertion Manager. This module manages the user insertion of content. The user can insert two types of content: *i*) she can insert new domain items (e.g., in the iCITY prototype, events), or details about existing items. This information is forwarded to the Domain Knowledge Manager, in order to be stored in the domain items knowledge base; *ii*) she can insert tags, comments, and ratings about existing domain items. This information is forwarded to the Annotation Manager in order to be stored in the Folksonomies (Annotations) knowledge base. The Content Insertion Manager supports the user insertion process, and in particular her tagging activity, by suggesting tags referring to the category in focus (if any) or suggesting the most popular tags. These suggestions can be personalized, i.e. based on user preferences, provided by the User Modeling Manager.

Moreover, the Content Insertion Manager acts as a “filter” on the inserted content. For instance, when the user inserts free text (as in an

¹ XHTML Mobile Profile is a subset of XHTML for WAP 2.0.

item description, or in a comment) it checks if such a text is acceptable with respect to a “black list” of forbidden words.

Finally, the Content Insertion Manager is in charge of extracting tags from item descriptions by exploiting the domain knowledge (DOM Ontology) that provides the domain vocabulary.

Dialog Manager. The Dialog Manager handles the interaction with the user. It has not reasoning capabilities. The Dialog Manager receives the web page from the Presentation Adapter and displays it. Moreover: i) forwards to the Content Insertion Manager all the contents inserted by the users (see above); ii) forwards to the Recommender the user requests concerning the presentation of (personalized) information; iii) forwards to the Context Manager context-related data, such as the type of device used by the user to interact with the system, the user location, etc.; iv) forwards to the User Modeling Manager user-related information - i.e. the data entered in the registration form; v) logs all the information concerning user actions (link selections, bookmarking, insertion of tags or comments, etc.) in a log file (see LOG box in Figure 1). Such logs will be accessed by the Behavior-Tracking sub-module of the User Modeling Manager (see dotted line between LOG box and User Modeling Manager in Figure 1) in order to learn new knowledge about the user, as described above.

3. A Web 2.0 enhanced user model

Web 2.0 gives adaptive web-based systems a great chance to enrich and extend the user model. Tags can be useful in increasing and optimizing the system knowledge about users (Carmagnola et al., 2007). In order to consider this *enhanced user model* in PAPERRE, we propose to analyze social annotations, and primarily, the tagging activity from a quantitative and a qualitative point of view. The quantitative point of view considers the *action of tagging* and of *annotating resources*² *per-se*, while the qualitative point of view reasons about the *meaning of tags*.

We should notice that the quantitative point of view refers to all the participative actions (tagging, commenting, annotating, and inserting new contents), while the qualitative point of view only considers the tagging activity. This choice is justified by the fact that the analysis of free-text comments requires a natural language component, which is currently not included within the framework. Instead, the analysis

² “Annotating resources” here means inserting those annotations different from tags, e.g. comments.

of tags can be achieved also without such a component, as it will be explained in the following sections.

3.1. ANALYZING THE ACTION OF TAGGING/ANNOTATING

According to the quantitative point of view, tagging, and more generally annotating, are considered as possible actions a user can perform on a social web site. Like other kinds of *usage data* (Kobsa et al., 2001) (clicking, buying, etc.), these actions represent an important feedback from the user. In fact, a user usually tags in order to highlight and organize the items she is interested in, so she can retrieve them later. Thus, the action of tagging is a stronger indicator for user interests (Kobsa et al., 2001) than simply clicking on a link, and therefore should be analyzed in order to make interesting inferences on the user model.

Table I offers an overview of this Web 2.0 *enhanced user model*. It presents a set of user model dimensions (columns from 3 to 11) that can be inferred from the action of tagging (first row), and from the meaning of the inserted tags (other rows)³. The categorization of tags has been derived from a user experiment that will be described in detail in Section 5.

In the following, all the dimensions presented in Table I will be detailed. For a quick interpretation of the table, notice that for each feedback and for a set of user modeling dimensions, the table specifies if the dimension is influenced by the feedback and provides the strength of such a correlation. For instance, if we consider the first feedback, i.e. “action of tagging”, the table reports a strong impact on three user model dimensions, i.e. “Interest”, “Interactivity Level”, and “Organization”. Notice that the correlations in the table have been assigned during a preliminary knowledge-gathering process, by interviewing user modeling and HCI experts of our research group. However, in principle, they could be revised or refined if needed. Let us now analyze in more detail how each user model dimension can be affected by the action of tagging.

³ Notice that almost all the types of feedback shown in the table are available thanks to the adoption of a Web 2.0 approach and would not be available in traditional adaptive systems. This shows that annotation provides important leverage on refining and acquiring information for the user model.

Table I. Overview of the Web 2.0 enhanced user model of PAPERE

Feedback	Interest	Interactivity Level ^g	Organization	Participation ^b	Knowledge ^c	Trust ^d	Creativity ^e	Abstraction ^f	Practical knowledge ^g
<i>Action of tagging</i>	Strong	Strong	Strong						
Category of tag									
<i>Proposed tags</i>			Weak		Weak				
<i>Most popular t.</i>			Weak		Weak				
<i>User's tags</i>			Strong		Weak				
<i>Recommended t.</i>			Weak		Weak	Medium			
<i>Free tags</i>			Medium		Medium		Medium		
<i>Known tags</i>			Medium		Medium		Medium		
<i>Unknown tags</i>			Medium		Weak		Medium		
<i>Specific tags</i>					Medium-strong				
<i>General tags</i>							Medium-strong		
<i>Synonym tags</i>					Medium-strong		Medium-strong		
<i>Spatial tags</i>									Medium-strong

^a The Interactivity level dimension, namely the measure of how much the user interacts with the system, considers not only insertion of new contents, but also the whole number of user/system interactions, such as clicking.

^b The Participation dimension considers how much the user is active and inserts new contents in the community.

^c The Knowledge dimension measures how much a user is supposed to know about a subject.

^d The Trust dimension expresses how much the user follows the system's suggestions.

^e The Creativity dimension expresses how much the user has the power to create new contents.

^f The Abstraction dimension expresses how much the user is skilled in classifying contents.

^g The Practical knowledge dimension considers the user knowledge about spatial and temporal features of items.

User interest represents how much the user is interested in the contents she annotated. If a user spends time in inserting tags or comments on a specific item, she is probably interested in that item. Traditionally, if a user clicks a specific item which belongs to a class of contents Cl , the system may infer an interest of the user in the class Cl with a low strength. The low strength is due to the fact that users may browse an item for many reasons (e.g. simple curiosity), not only for a real interest in it. A stronger correlation may be assumed if the user also performs other actions on resources of the class Cl . In our model, if a user adds some tags to a specific item which belongs to a class Cl , this action can be a sign of interest in the item, and thus in the class Cl , with a medium-high degree of strength. It is, in fact, quite unlikely that a user would spend time tagging an item, or writing comments on it, if she is not really interested in the category it belongs to. Thus, the system inference about user interests can be more accurate and quicker than the traditional user modeling process.

User interactivity level is the measure of how much the user interacts with the system. It is related, on the one hand, with the willingness of the user to interact with the application (how much interactive a user is), and, on the other hand, with the concrete possibility for the user to interact with the application (how the contextual situations encourage the user's interactivity). The action of tagging seems to be a relevant indicator of the user interactivity level, since performing it requires some effort and thus the more a user adds tags to some resource, the more she can be considered as an active user. Likewise, other participative actions enabled by the instantiation of PAPERÉ (see Section 4.5), such as annotating resources or inserting new events, can reinforce the system assumptions about the user interactivity level.

User organization level identifies the attitude of the user in organizing and categorizing things. The main motivation for the user to tag on the web is to satisfy her need to store and organize resources in a personal way in order to better visualize and retrieve them later. The system could analyze the user cross-tagging activity, namely how often the same tag is used to categorize different resources. From this analysis, the system could infer that, if the same tag is frequently used to annotate different items, it is probably used to group items sharing similar features. Thus, this way of using tags is strongly related to the organization level of the user, as specified in Table I.

3.2. ANALYZING THE MEANING OF TAGS

In Section 3.1, we have discussed the potential inferences based on a quantitative analysis of tagging. The qualitative point of view investigates the possibility to reason on the meaning of tags. Tag semantics could be analyzed by means of: *i*) an ontology of the application domain, including relevant concepts and the relationship of the concepts and *ii*) a database that specifies a list of terms for each concept in the ontology. The latter can be either a controlled vocabulary⁴ (as in the iCITY prototype) or a lexical database, such as WordNet (Fellbaum, 1998).

However, it is interesting to investigate the meaning the tags may assume in relation to their specific context of use. This is the reason why our classification takes into account the experiment of Carmagnola et al. (2007), described in detail in Section 5. This work has identified a number of categories of tags related to their usage (Proposed Tags, Free Tags, General tags, Specific Tags, Synonym Tags, Contextual tags, and Subjective Tags). This classification is used in the user modeling knowledge base of the framework, in conjunction with the domain ontology and the lexical database, in order to enrich the system interpretation of tags, as detailed below. Notice that, with respect to this classification, we do not consider *Subjective Tags* - namely tags that express user opinions and emotions (e.g. *boring*, *interesting*, etc) - since we do not have, at the moment, a comprehensive classification of adjectives that can help at clarifying their positive/negative meaning. However, iCITY is able to understand a user's degree of satisfaction with an item. In fact, users have the chance to rate each event explicitly by choosing the number of stars representing their satisfaction; from this action, the system can easily infer the users' positive or negative judgment. Below, we describe each category of the tag classification, which user model dimensions each category may have an impact on, and the degree of correlation between the feedback and the impacted dimension.

Proposed tags/free tags. *Proposed tags* are tags proposed by the system. They are further categorized in: *most popular tags*, the tags most used by the community of users; *user's tags*, the tags mostly used by the user; *recommended tags*, the tags deriving from the system inference. The user may prefer to select *proposed tags* for various reasons: she may be a bit lazy; she may not have time to insert her own tags; she may not have a deep knowledge on a topic and so try to select tags from words presented by the system; she may not be so interested

⁴ A controlled vocabulary is any subset of the natural language lexicon which is organized as a list of preferred terms, often called *authority file*.

in personally categorizing the items. Notice that the strength of these inferences is weak, since the user behavior could be due either to the mentioned reasons or simply to the fact that the user has found the right tag among those suggested by the system. In particular, considering the subcategories, if the user selects some *most popular tag*, we could weakly infer that she trusts the other people of the community and that she conforms to the general thought (conformism). However, if she always uses the same tags after some interactions (*user's tags*), we could infer that she very probably has a high propensity for regular habits and organization. Finally, if the user selects a *recommended tag*, we could infer that the user trusts the system (with a medium strength that can increase during subsequent interactions).

Free tags are tags the user freely enters. They can belong to the system vocabulary (*known tags*), if they correspond to some words in the controlled vocabulary or in the lexical database. Otherwise, they can be either invented (such as unhyphenated compound words like “picassoExhibition”) or simply unknown by the systems, in both cases classified as *unknown tags*.

The use of a lot of free tags can be directly related to: *i*) the knowledge in the topic, because inserting free tags may require a specific knowledge in the area⁵; *ii*) the level of creativity; *iii*) the level of participation in the tagging activity, because using personal words requires more effort than simply selecting them from a list of suggested tags; and *iv*) the level of organization, because a user may intend to annotate resources personally.

General/specific tags. *General tags* classify the resource in a general dimension. Tags may be considered as “general” when it is possible to map them on an upper class of an ontology, or on hyperonyms of WordNet. *Specific tags* add some specification, such as specific words that can imply that the users have knowledge in the topic. However, the contrary (the use of general words) does not imply that users are lacking in knowledge. In fact, if the general words are appropriate, this could imply that the users have good abstraction capabilities.

Synonym tags. *Synonym tags* are words equivalent to the tags used to describe a resource. The new tag inserted by the user can be compared to a controlled vocabulary in order to find out if the term is a synonym of some words used to describe the specific item. For example, imagine a system that proposes an event which is classified as an instance of the domain ontology sub-class “jam session”. If the user tags the resource as “improvisation”, the controlled vocabulary might help the system to

⁵ This assumption can be reinforced by analyzing whether the free/known tag refers to the same concepts as the tagged resource.

classify the term as a synonym. The use of synonyms can be directly related to the knowledge in the topic and to the level of creativity (see Table I).

Contextual tags. *Contextual tags* are tags referring to the context of the resource. The use of contextual tags, specifically *spatial tags*, could mean that the user has a practical knowledge of the resource location, probably derived from a direct experience (e.g. visit of the tagged place, participation at the tagged event, etc.).

4. An application to cultural heritage: iCITY

iCITY⁶ is a web-based application which instantiates the PAPER framework. iCITY is an adaptive, social, multi-device recommender guide in the cultural heritage domain. It deals with cultural events taking place in Torino. The events are derived both from institutional sources (RSS feeds provided by the Torino Cultura portal⁷) and from informal knowledge shared by willing users. Thus, part of iCITY content is user-generated, in accordance with one of the most important Web 2.0 trends. User-generated contents do not only increase the amount of available information, but they also make the knowledge bases more diversified by integrating minor events, which are not usually covered by institutional sources.

In particular, registered users can post new events, ratings, textual comments, further details and tags. The idea is that allowing users to post their own contributions represents a useful means to refine contents, integrating both new knowledge and firsthand experience, thus enriching plain information with an evaluative dimension. Measuring themselves up against other people's opinions and experiences, users have more resources for making informed choices and for deciding whether a certain event is likely to interest them or not.

4.1. NAVIGATION

iCITY users can browse the events in both "text" and "map" mode (see Figure 2). In "text" mode, users are provided with a list of events, with a title and a brief description; a more detailed textual description can be accessed on demand. In "map" mode, events with available addresses are geolocalized and displayed on a map.

In addition, users can choose whether to navigate a taxonomy or a folksonomy when browsing events, both in text and map mode. In the

⁶ <http://www.icity.di.unito.it/dsa-dev/>

⁷ <http://www.torinocultura.it/>

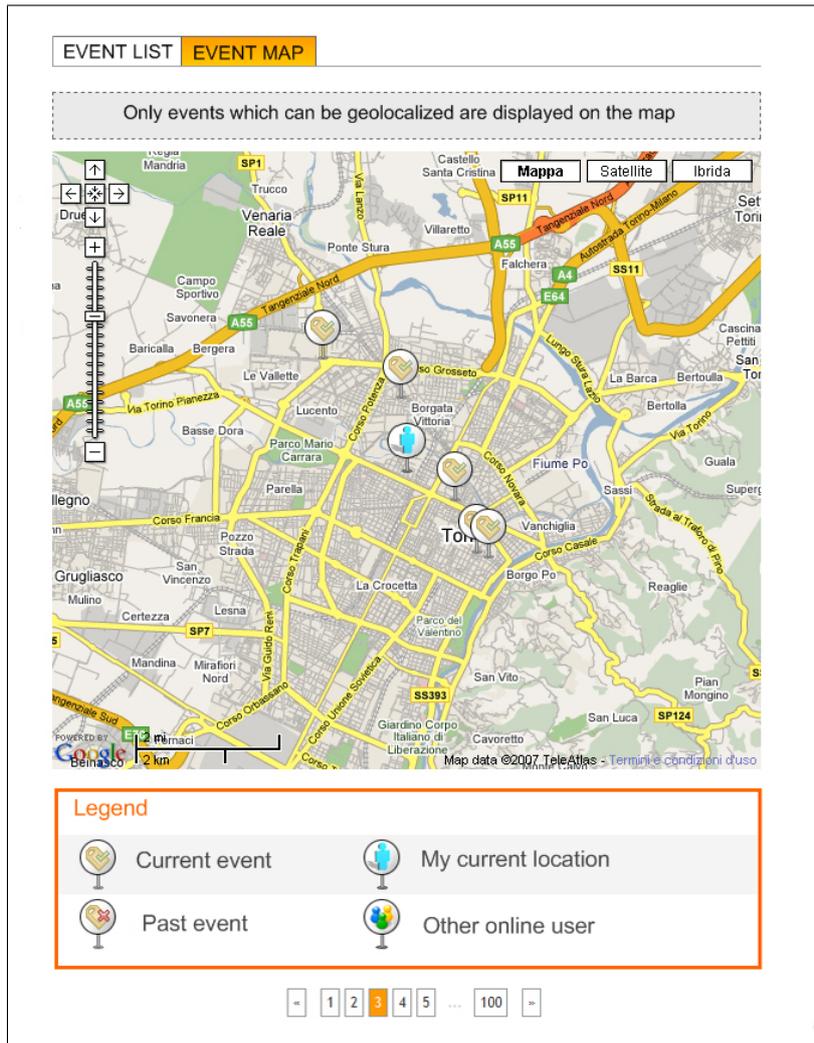


Figure 2. iCITY map: geolocated events are displayed on a map together with the current user location.

first case, they can filter the events by selecting one of the hierarchically organized categories or subcategories defined to classify the system content (see Figure 3). Users who prefer a “folksonomic” approach can use tag clouds for their navigation and view all the events which were labeled with a certain tag. Specifically, a user can browse her own cloud or the cloud containing the tags used by the community of users.

The taxonomic and folksonomic approaches are not mutually exclusive, but can be effectively interchanged at will in order to allow users

to enjoy both a top-down, well-structured exploration process and a bottom-up, dynamic one.

4.2. CONTENT PERSONALIZATION

Event lists are tailored to each user according to her interests, represented in the user model, and to contextual elements, e.g. her location and the events date. More specifically, events are ordered so that the most appropriate ones are displayed at the beginning. Recommended items are emphasized through the “adaptive annotation technique” (Brusilovsky, 1996), a navigation support that consists in attaching various visual cues (e.g. enlarging fonts, changing colors, adding icons, etc.) to the suggested items to help users to select the most relevant ones. In iCITY, recommended items are annotated with up to four colored thumbs-up icons, representing how much the item should interest a particular user (see Figure 3). Users are also provided with other information which should help them to quickly evaluate the events: the number of views and the rating. *Views* represent a quantitative index of the event “popularity”, while *rating* (represented by stars) is a mean value which summarizes users’ opinions about a certain event. Notice that, as demonstrated in (Cena et al., 2005), stars usually express a qualitative evaluation associated with an expert opinion or with a general assessment, and are often used in that meaning. Meanwhile, thumbs-up icons represent a more direct engagement with the user and communicate in a more direct way that “this item is ok for you”. Thus, thumbs-up icons visually represent how much the recommended items match the user’s preferences.

Users can influence the personalization process by indicating which criterion should be given more importance when the system selects the events to recommend. The available criteria are: *level of interest*, *rating*, *recentness*, and *proximity* of each event. The *level of interest* depends on how well a certain event matches the current user’s interests. *Rating* is the average of all the evaluations users have expressed about a certain event, while *recentness* and *proximity* relate to its date and location, respectively. Alternatively, users can indicate the criterion they are most interested in when they provide information for their profile. Users can change the score formula (see below) dynamically by selecting the most relevant criterion; in this case, the change will affect the current session only. When users select a different criterion, the event list is rearranged dynamically, so that users can enjoy a real-time feedback. Users can also select their current location or another location they are interested in. Location is particularly important since it directly impacts on map mode presentation. Maps do not only provide an immediate visual

The screenshot shows the iCITY website interface. At the top, there is a navigation bar with links for 'about us', 'contacts', 'help', 'FAQ', and a search icon. Below this is the iCITY logo with a 'BETA' badge and the text 'Digital Semantic Assistant TORINO'. A search bar is present with the placeholder 'insert a username' and a 'search' button. To the right, there is a user profile section for 'Proserpina' with options to 'logout', 'my location', 'my profile', 'my tags', 'my favourites', and 'my friends'. Below the profile is an 'inbox' notification for 7 messages, with 4 new ones.

The main content area is titled 'Events' and includes a 'You are here: home > all categories' breadcrumb. It features a 'recommend by:' section with links for 'interest', 'recency', 'user location', and 'rating'. On the left, there is a sidebar with 'All categories' and a list of event categories: Dates, Art, Cinema, Books, Music, and Theater. The main event list shows three items:

- The Painting of Modern Life**: Inserted 8 days ago, from Torino Cultura. Event classified in Art > Museums. When: from 06 Feb 2008 to 04 Mag 2008. Where: Piazza Marfaldia di Savoia. 0 views.
- Senti le lingue?**: Inserted 8 hours ago, from Torino Cultura. Event classified in Dates > Meetings - Conferences. When: 14 Feb 2008. Where: Corso Orbassano. 0 views.
- Tesori nascosti dalla terra dei faraoni**: Inserted 2 months ago, from Torino Cultura. Event classified in Dates > Meetings - Conferences. When: from 11 Ott 2007 to 14 Feb 2008. Where: Corso Stati Uniti. 3 views.

On the right side, there are two tag clouds: 'TAGS I USE THE MOST' (No tags are available) and 'THE MOST POPULAR TAGS' (listing tags like art, cinema, movie, meetings, reading, books, exhibitions, museums, music, orchestra, presentations, show, theater). Below these is a 'SIMILAR USERS' section listing 'max fra estragon' and 'chicco josephine'.

Figure 3. The list of cultural events is personalized according to four criteria: user interests, user location, event recentness and event rating. Recommended items are annotated with up to four thumbs-up icons. Two tag clouds with the most popular tags are displayed on the right

feedback for the effectiveness of personalization but they also allow users to experience contents in a more immediate way.

iCITY users can inspect the system assumptions about their interests by accessing their *scrutable user model* (Kay, 2006). More specifically, users are provided with a table summarizing their levels of interest in each category of the event taxonomy, as inferred by the system. At the moment, users are not allowed to edit their user model explicitly - e.g. by providing weights or correcting the system assumptions. Nevertheless, they can help better tune it by providing some tags describing their interests in their profile page.

4.3. PERSONALIZATION AND COMMUNITY

Registered users have the opportunity to benefit from iCITY community. In fact, they can create their own social network by defining other users as “friends”. Users can view their friends’ current location on a map when they are online and they are also notified if their friends

significantly interact with the system - e.g. if they post new events or comments.

Furthermore, each user is provided with a list of similar users that can be used as a different starting point for exploring the resources in the system. A similarity value is automatically determined by the system for each possible pair of users by confronting their level of interest for the various categories and subcategories. When such a degree of similarity crosses a certain threshold, the two users are considered as similar - i.e. they are considered as having common interests and preferences. By accessing the “profile page” of similar users, the current user can learn some more information about them (e.g., age, gender, short description) and analyze the list of tags they used to describe and classify interesting events they browsed in iCITY. Personal information is only displayed if users agree to share it, by managing privacy settings in their profile page. Consequently, the current user is offered another way to easily retrieve potentially interesting events, basing on the assumption that she shares some interests with similar users. In order to provide the current user with more opportunities, it might be useful to enrich the similar users’ profiles with a list of the events they posted.

4.4. ADAPTIVE PRESENTATION

Whereas contents are mainly tailored according to the user interests, the presentation is adapted on the basis of the device, distinguishing between *desktop* and *mobile* devices. We designed two *mobile* interfaces. There is a simple interface, suitable for mobile phones and, in general, for the so called Default Delivery Context, as defined by the W3C⁸ and there is a slightly enhanced interface, which can be accessed by PDAs. Both interfaces enable users to perform the same tasks, but the PDA version features a richer layout.

Content presentation and navigation structures are arranged in order to fit mobile users’ needs (see Figure 4). For example, the number of events per page was reduced to two, while the three-column layout of the desktop version was abandoned in favour of a one-column layout, which only requires vertical scrolling. In addition, internal anchors were provided in order to facilitate page navigation, at the same time minimizing scrolling. Due to display limitations, navigation structures are arranged in a more “goal-oriented” perspective: for example, taxonomic and folksonomic navigation were grouped together, since they both enable users to browse events. On the contrary, links to users’ personal pages are grouped in a different section. In order to provide an adequate

⁸ <http://www.w3.org/TR/mobile-bp/>

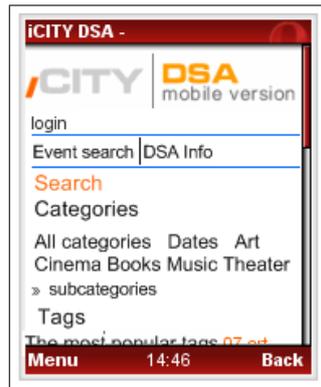


Figure 4. iCITY version for mobile phone.

user interface for each of the contexts and device categories, we designed a clear separation between content and presentation. In designing both desktop and mobile interfaces, usability and accessibility issues were addressed. The whole application was designed specifically for usability through applying a user-centered methodology that implied a series of heuristic evaluations, usability tests and consequent redesigns (see Section 5).

4.5. iCITY ARCHITECTURE

In Section 2 we presented the architecture of the PAPER framework. Now we will describe how the general architecture has been instantiated in the iCITY prototype (Figure 5). Many of the architectural possibilities offered by the framework have been simplified in order to fulfill the domain requirements.

4.5.1. Knowledge representation

iCITY, with respect to the PAPER framework, employs lightweight ontologies in the form of taxonomies.

Domain KB

The knowledge base that stores the knowledge of the system about the cultural domain is represented by means of a *taxonomy*. Although this choice has reduced the expressive power of the representation, taxonomies are more efficient in reasoning mechanisms. The classes in the domain taxonomy of iCITY represent event categories. Such a structure is derived from the information source, Torino Cultura, which provides events classified in a taxonomy of channels and sub-channels. The Domain KB of iCITY also includes a *controlled vocabulary* (or *authority file*). Words in the controlled vocabulary and elements in the

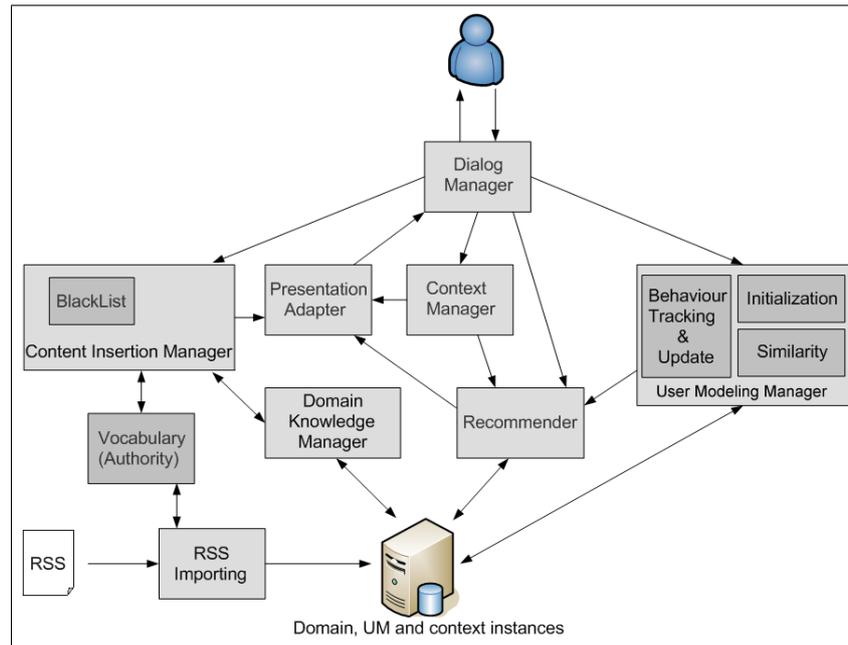


Figure 5. Architecture of the iCITY prototype.

events taxonomy have a many-to-many relationship. This relationship enables the system (specifically, the Content Insertion Manager, see below) to support the final user actively, for instance, by suggesting concepts from the vocabulary when the user is tagging an event. Moreover, thanks to this relationship, when a user inserts a tag among those suggested by the system and belonging to the controlled vocabulary, the system is able to map it to the taxonomy and thus derive the meaning of the tag.

Folksonomies (Annotations) KB

Tags can be “free” - i.e., inserted by the user - or “controlled” - taken from the dictionary associated with the Domain KB. *Comments* can be *ratings* - provided by the user to express her opinion about the event on a Likert scale - or *textual evaluations* expressing a critical judgment. Currently, the system only reasons on the action of tagging (see Section 3) and on ratings⁹. The action of tagging will be described below; concerning ratings, iCITY combines all the ratings on an event and calculates an average value expressing the global judgment of the users about such an event.

⁹ In the current version, the system is not able to process natural language and thus textual comments are just displayed.

Domain items and Info Provider: RSS Importing

To automatize the uploading of up-to-date information about cultural events, the system uses a stream of RSS feeds provided by Torino Cultura. The framework can read several different formats beside RSS feeds. The only requirement is that the imported content must be suitable to be classified into the domain taxonomy. In other words, it must be possible to automatically map the imported contents to the elements of the domain taxonomy. When contents are imported, they are transformed from a flat representation (a string) into a structured object by means of parsing mechanisms based on finite state automata. The reference pattern is to use an *adapter* that performs the adaptation of the imported format into the structure used by iCITY to represent items. Thus, to be able to import events from heterogeneous sources, the system is simply required to write the right adapter.

Context KB

In iCITY, the modeling of the context includes two main aspects:

- The user device. iCITY automatically detects the *typology* of device, by referring to a taxonomy of devices, and tailors the presentation of contents to it (see the Personalization Adapter described below).
- The geo-localization of events and other users. This is a crucial functionality for iCITY, because it allows the user to know *where* an event takes place. The geo-localization is expressed by means of GPS coordinates and is represented through interactive maps¹⁰. The application translates textual labels into GPS coordinates using the Google Geocoder service¹¹.

Moreover, although the system does not include any ontology to model time, it is able to deal with the information regarding the temporal validity of events. Thus, it can satisfy user needs regarding *when* an event takes place.

User Model (UM) KB

The user model in iCITY maintains personal information and a multi-dimensional representation of the user interests in each category of the domain taxonomy, by associating to each category a probability value.

¹⁰ In particular, iCITY exploits the Google Maps available as free APIs: <http://maps.google.com>

¹¹ <http://www.google.com/apis/maps/documentation/services.html#Geocoding>

Thus the user model contains a probability distribution of user interests. Table II shows an example of a User Model, with values associated to user interests, after a number of updates.

Table II. Example of user model's feature-value pairs

User model's feature	Value
Art - Architecture/Design	0.023181
Art - Multimedia/Performances	0.013145
Art - Museums	0.001416
Cinema - Movies	0.004132
Cinema - Documentary	0.008173
Books - Readings	0.012316
Music - Rock	0.023189
Music - World Music	0.056131
Music - Jazz/Blues	0.032311

Notice that the model is initialized to a uniform probability distribution representing the same interest value for each domain category. Such values will be modified in the further update phases.

Moreover, the user model contains information about the tags and comments inserted by a user, which represent her subjective view of the system contents.

4.5.2. System modules

In this section we describe the function of the system modules, with a particular attention to content adaptation which represents the most relevant “intelligent” behavior of the system.

Dialog Manager

The Dialog Manager is the module that is in charge of the interaction between the system and the user. It is implemented as a Java Servlet that receives the user requests and replies to them, and represents the only input/output point. This allows a complete and secure monitoring of data exchange between a user and the system. The Dialog Manager logs all the user actions into a log file which will be accessed by the User Modeling Manager (see below). Moreover, the Dialog Manager dis-

patches the different user requests to the modules in charge of handling them - i.e., the Recommender and the Content Insertion Manager. Such modules, in turn, send their responses to the Presentation Adapter, which tailors the presentation to the user device (see below). Finally, the output of the Presentation Adapter is forwarded to the user by the Dialog Manager.

Presentation Adapter

This component adapts the user interface to her device. In order to identify the device type, the system detects the user-agent (extracted from the HTTP request header) and matches it to a class in the device taxonomy.

For each user request, the contents selected by the system are represented as XML documents. The Presentation Adapter selects the proper XSL stylesheet in order to perform a transformation that tailors the user interface to the device characteristics by producing XHTML documents for desktop (and laptop) computers and XHTML MP documents for mobile devices¹².

The selection of the XSL stylesheet is performed by means of rules like the following:

IF (device=PDA) THEN (stylesheet=medXSL.xsl)

where the left hand side of the rule matches the user-agent (*device*) to classes in the device taxonomy (namely, *PDA*) and the right hand side provides the *stylesheet* to be applied to perform the proper transformation (*medXSL.xsl*).

Content Insertion Manager

This component handles the information inserted by the users. For each request containing information to be inserted into the system knowledge bases, the Content Insertion Manager (CIM) checks the content using some filters - e.g. to avoid words in a *blacklist*. It also analyzes the inserted tags according to the vocabulary (as described above). Finally, it supports the user tagging activity by suggesting tags that refer to the category in focus or the most popular categories.

Recommender

The recommendation of contents is based on different criteria related to user's features and behavior, and to event characteristics:

- *Interest*: the level of user's interest for each subcategory of events (*intScore* in formula (1)).

¹² At the moment, the system supports desktop and mobile platforms (PDAs and Smartphones). We are developing a version for Web TV (accessible through the Media Center).

- *Position*: the current position of the user ($(posScore)$ in (1)). This value is normalized considering different levels of proximities, as in (Cena et al., 2006).
- *Recentness*: the temporal delay between the current date and the date of the end of the event ($(recScore)$ in (1)); the smaller the delay, the closer the event, thus this value is complemented to 1.
- *Rating*: the global judgment calculated according to user’s comments on a particular event ($(ratingScore)$ in (1)). The value is normalized to 1.

According to these criteria, for each event, a total *score* is computed using the following formula:

$$eventScore = intScore * w_1 + posScore * w_2 + recScore * w_3 + ratingScore * w_4 \quad (1)$$

where w_i represents the weights associated to each of the four above-mentioned criteria. An example of weight distribution which favour user interests and location is the following:

$$w_1 = 0.5 \quad w_2 = 0.3 \quad w_3 = 0.15 \quad w_4 = 0.05$$

All the partial scores are normalized as a value in the interval $[0, 1]$. Once the score ($eventScore$) for all the events related to the user’s request are computed, the events are ordered on a descending scale. The events with the highest score are what the system considers the most interesting for the user and so they are presented at the beginning of the list.

User Modeling Manager

The User Modeling Manager initializes the User Model, updates it, and calculates the similarity among users by means of the following sub-modules.

i) The Initialization sub-module. It initializes a new User Model by inserting data coming from the user registration form - e.g. nickname, etc. - and by assigning a uniform probability distribution to the user interests for domain categories (types of events: music, theatre, and so on).

ii) The Behavior-Tracking sub-module. The update of the User Model is one of the most crucial tasks for the system. It is accomplished by considering the actions performed by the user when interacting with the system¹³.

¹³ Notice that in the current prototype we exploited tags only to perform a quantitative analysis, thus considering the action of tagging (see Section 3.1). The

Table III. Weights associated to actions

Action	Weight
Adding events or bookmarking	0.9
Tagging	0.8
Updating information about an event	0.7
Inserting a comment/rating	0.6
Visualizing an event or clicking on a detail on the map	0.3

The different types of actions tracked by the system are weighed in a different way as they provide different pieces of evidence about the actual user interest. The weights have been assigned according to the ideas expressed by Kobsa et al. (2001) and our past experience with UbiquiTO (Cena et al., 2006). The weights we are currently using are reported in Table III.

The Behavior-Tracking sub-module, for each user interest in the User Model, applies the following formula:

$$computedValue = \frac{\sum_{i=1}^{actions} (count(i) * actionWeight)}{totalWeight} \quad (2)$$

For each action type i , $count(i)$ is the total number of actions of type i performed by the user and $actionWeight$ is the weight of the action type i , according to Table III. The formula sums up the results of the product of $count(i)$ and $actionWeight$, calculated for each action type i ; i ranges from 1 to $actions$ that is the number of action types the system is able to detect. The total sum is then divided by $totalWeight$ which represents the sum of the weights of all action types. Thus, for each interest in the User Model, $computedValue$ represents the value computed according to user behavior (i.e. based on user actions). By applying this rule to all the user interests in her User Model, the Behavior-Tracking sub-module obtains a new probability distribution for all user interests.

This new probability distribution is then combined with the one in the current user model to update it. For each item in the probability

qualitative analysis of tags, namely the meaning of tags (Section 3.2) is planned for the second version of iCITY.

distribution, we perform the update as follows:

$$newValue = computedValue * w_1 + currValue * w_2 \quad (3)$$

where *computedValue* represents the probability value computed on the basis of the user actions, *currValue* is the probability value in the User Model to be updated, and w_1 and w_2 (summing to 1) are used in order to give different weights to the two contributions. At the beginning of the interaction, w_1 is close to 1 and w_2 is close to 0 (the initial User Model is not significant because it is the same flat model for all users). Then w_2 is progressively increased and w_1 is decreased. This is done in order to assign increasing weight to the current user model, as long as it begins to represent a quite long history of interactions.

iii) Similarity submodule. The task of calculating the similarity among users is important, since iCITY aims at being a *social software* - i.e. a tool for creating and managing a community of users. As borne out extensively in the literature, see for instance (Schafer et al., 2007), the formula for computing the similarity among users is based on a variation of the formula for the *standard deviation*, where the values of the user interests are used instead of the expected values (mean values). The formula can be expressed as follows:

$$similarity(user_1, user_2) = \sqrt{\frac{\sum_{i=1}^{interests} (user_1[i] - user_2[i])^2}{interests}} \quad (4)$$

where $user_1$ and $user_2$ are two generic users, i represents each interest and *interests* is the number of interests. For each couple of users, the system compares the value of corresponding interests. In order for one user to be considered similar to another, their similarity must pass through a threshold (currently 0.7).

4.6. IMPLEMENTATION REMARKS

iCITY is mainly based on open-source technologies, conforming to the most important standards.

We implemented our system by using the Java Servlet technology in an Apache Tomcat environment. The inference engine embedded by the Presentation Adapter is based on Jess (Java Expert System Shell, a Java library). User and event data are stored into a MySQL database. To visualize the localization of events and users, iCITY uses Google Maps, available through public APIs.

For the dynamic generation of the different user interfaces we produce an XML document which represents the content of the interaction. Then we use XSLT to transform this document into XHTML

or XHTML Mobile Profile, depending on the device. All the results respect W3C standards¹⁴ (XHTML Transitional, CSS 2).

Moreover, iCITY follows W3C recommendations regarding accessibility and usability of the system.

Finally, iCITY is able to import and export events represented with a RSS 2.0 format¹⁵.

5. Evaluation

As proposed in Gena and Weibelzahl 2007, we carried out different evaluations at different stages of development: the requirement phase, the preliminary evaluation phase and the final evaluation phase.

The requirement phase. The requirement phase is usually the first phase in the system design process. During the initial phase of early design of the framework implementation, we carried out an evaluation in order to gather requirements for exploring how users tag items. The objective was to design the user-modeling knowledge-base of iCITY, especially collecting feedback from tags and annotations. Therefore we were interested in discovering how real users use tags.

To simulate the tagging activity, we chose a list of events from the RSS channel that feeds iCITY. We selected 15 events belonging to different categories - art, theater, cinema, music and books. Then we grouped them and presented them to three different groups of users. We selected 39 users among students (23 subjects), researchers working in our departments (10 subjects) and relatives and friends (6 subjects). Notice that subjects were selected according to an *availability sampling strategy*¹⁶.

The test was carried out as follows. We presented each user a printed list of 5 *events* with their corresponding *description* and we asked them to tag the events, as if they were interacting with a social software. Users were free to write their own tags (up to 5 tags for event) or to choose them from the description of the resources. This second option aims at simulating the suggestion of tags a system can support. Notice that the testing design was aimed at simulating the real behavior of

¹⁴ <http://www.w3.org/>

¹⁵ <http://www.rssboard.org>.

¹⁶ Even though random sampling is the best way of having a representative sample, these strategies require a great deal of time and money. Therefore much research in psychology is based on samples obtained through non-random selection, such as the *availability sampling* - i.e. a sampling of convenience, based on subjects available to the researcher, often used when the population source is not completely defined (Royce and Straits, 1999).

iCITY, which presents events with their description, initially annotated with tags extracted from the event description. At the end, we collected 217 tags and we analyzed them inductively, according to the principles of the Grounded Theory (Strauss and Corbin, 1990).

In particular, for our aims, we focused our analysis on the first two phases the Grounded Theory proposes. Phase 1, *open coding*, is the inductive process through which concepts are identified starting from the analysis of the collected material, and then grouped into categories. Phase 2, *axial coding*, is the process of relating categories to their subcategories, linking categories at the level of properties and dimensions.

We examined the tags, and we compared them for similarities and differences, starting to accumulate concepts. At the same time we began the inductive process of investigating and defining the main categories, subcategories and involved variables. The first two categories we considered, and their corresponding frequencies in the test, are: *i*) proposed tags (tags derived from the resource description): 76% of the tags; *ii*) free tags: 24% of the tags.

Taking into account other properties related to the tagged resource, the following sub-categories emerged. There were *specific tags* (tags that add some specification about the resource): 61.19% of the tags. There were *generic tags* (tags that classify the resource in a more general way): 22.37% of the tags. There were contextual tags (tags about the context of the resource: location, time, etc.): 13.24% of the tags. There were *synonym tags* (tags that are synonyms of terms in the resource description): 2.74% of the tags. There were *unknown tags* (e.g. unhyphenated compound words like “picassoExhibition”): 2.17% of the tags.

Notice that these categories are similar to the ones identified by the group which is working on “My Web 2.0”, the social search engine of Yahoo! (Xu et al., 2006)¹⁷. In contrast to that analysis, our classification lacks two categories: *subjective tags* (tags that express user’s opinions and emotions) and *organizational tags* (tags that identify personal items). We did not find any occurrence of these kinds of tags. A reason can be that taggers typically use subjective tags with social and organizational aims. Both kinds of tags are meaningless in a test environment and thus our sample users did not feel the need to use them.

Considering this gap, we integrated the classification with the categories that did not emerge from the test, i.e. *subjective tags* and *organizational tags*. Notice that, in iCITY, free comments and ratings

¹⁷ For another classification see Golder and Huberman (2006).

the users can post in a dedicated area, are considered as forms of subjective tags since the user is able to express her qualitative and quantitative judgment about events. Moreover, we also integrated the classification with the typologies of tags social softwares frequently suggest : *i) most popular tags*, the most popular tags in the community; *ii) most used tags*, tags previously inserted by the user; *iii) recommended tags*, tags provided by the system.

Since the results that emerged from this classification seem to be general enough, we decided to include the classification in the PAPER framework too, as described in Section 3.

The preliminary evaluation phase. The preliminary evaluation phase occurs during the system development. During the development of iCITY we carried out two different evaluations: a heuristic evaluation and two sets of usability tests.

The *heuristic evaluation* was carried out by a HCI expert and an adaptive web expert, who ran their heuristic studies separately and conducted independent evaluations. These evaluations were aimed at checking *i) the usability of the interface and the conformance to general HCI principles*; *ii) the usability of the adaptive behavior of the system, based on the existing literature in adaptive systems evaluation* (Jameson, 2006; Gena and Weibelzahl, 2007). In particular, the experts were asked to follow Jameson's 5 usability challenges (Jameson, 2006) for adaptive interfaces. This evaluation led to a first re-design both of the user interface and of some aspects of the system functions. In particular, given the principles proposed by Jameson, the two experts found the following problems. First, *predictability* and *transparency*: to solve these problems, the adaptive behavior was made explicit through labels explaining the filtering modality of recommendations and the user model was made scrutable. Second, *controllability*: the user was given the chance to sort recommendations according to different parameters (e.g., interest, location, rate, etc.). Third, *unobtrusiveness*: the user was given the chance to add information on her user profile, but this was not mandatory. Fourth, *privacy*: the scrutable user model and the bookmarked events were made unavailable to other users. Fifth, *breadth of experience*: the less suggested events were positioned in a low ranking position, instead of being filtered out. Moreover, the ranking was made modifiable by the user.

Then we carried out *two sets of usability tests of the scenario-based prototype*. More specifically, we designed two different sets of tasks in order to guarantee that all the features were evaluated. We tested each set with five users. These users were selected among target users of the system and were between 25 and 40 years old, both female and

male. All of them were familiar with the web, but not all of them were accustomed to typical Web 2.0 features. Each test session took about 20 minutes and was recorded for later reference. The usability problems we discovered, thanks to these tests, were addressed in the second re-design of the system interface.

In particular, usability tests showed that some labels were ambiguous, so that users were unable to match them with the corresponding functions. For example, the label “add an event” was perceived as equally likely to refer to the functions “post a new event” and “add an existing event to your favorites”. Users also complained about the lack of immediate feedback on some actions, such as updating their profile or adding a new contact to their list of friends. Long descriptions in the event list were considered redundant and not very useful, since they partially overlapped with the detailed event descriptions. Moreover, users felt overwhelmed by big tag clouds with many similar-sized items and needed detailed instructions in order to cope with the several similar forms. The map-based modality of interaction was also novel and quite puzzling for some users, thus requiring a few hints. We tried to address all these problems, which were solved in the subsequent re-design.

The final evaluation phase. The final evaluation phase occurs at the end of the system development and it is aimed at evaluating the overall quality of a system with users performing real tasks. Thus we decided to test the system under real conditions with real users performing real tasks. Users were also asked to report any kind of problems they experienced while using the system.

We selected a group of 20 users, 23-40 years old, 12 females and 8 males, among colleagues and friends, with the following features. All users were target users of the system: skilled computer users, frequent Internet users, familiar with technology, living in the city of Turin and really interested in cultural events taking place in the city. Notice that subjects were selected according to an *availability sampling strategy*. We asked them to register onto the system and to use the “desktop version” of the system every day for two weeks. We did not give any instruction, since we did not want to influence the testers with our suggestions.

After this period, users compiled a free report containing the problems they experienced. Then they were asked to evaluate their scrutable user model, containing the scores the system assigned to each *category of interest*. These scores were normalized to sum up to 100. Thus, each value was represented as a percentage. Users were asked to evaluate if scores matched their real interest in each event category.

If not, they had to assign the correct values and send them back to the experimenters by email. 8 users out of 20 said the user model was correct, while 12 subjects re-assigned part of the values, thus expressing different interests with respect to the system inferences¹⁸. In order to evaluate the accuracy of the system predictions, we calculated the MAE (Mean Absolute Error), a statistic metric that estimates the accuracy of a filtering system by comparing the numerical values of the system predictions against user ratings for the items that have both predictions and ratings (Good et al., 1999). In particular the MAE measure evaluates the distance between the system predictions and the user's opinion using rate vectors. The smaller the MAE value, the greater the prediction accuracy. We obtained a medium MAE of 0.11, in a range from 0 to 1, which in literature can be considered as a good value¹⁹.

We also carried out unstructured interviews with the subjects, which were very fruitful and led to a third re-design of the system. We asked them opinions about the system: what they liked, what they did not like, what needed improvements, and whether they appreciated the recommendations. Their free answers were deeply investigated by the interviewer in order to gather suggestions for an improvement of the system as well as to collect feedback on the interface, the tag-based approach and the recommendations. Apart from usability problems and minor bugs, users were quite satisfied with the system. In particular, the 8 users whose user model was correct referred that the system perfectly reflected their interests in that period. The remaining subjects, instead, criticized the fact that some interests were overestimated. This can be explained by the fact that the initial user model changes rapidly (see Section 4.5.2), on the basis of the user's actions and thus, in a limited period of time as in the test, it is liable not to conform with her interests. Thus we enlarged this period. Moreover, we investigated how these users use the system by inspecting what kind of actions they performed in the real context of use. Users do not only appreciate the chance to be informed about the events promoting cultural heritage of the city, but they also liked to be able to add new events in order to exchange information with other users. For instance, we noticed that users tend to use the free comment form as a place for asking information to other users of the community (e.g., "do you know if tickets are still available for that concert?"). As an improvement, users requested a utility to group the events they inserted. Moreover, they expressed appreciation for the different ranking modalities, when look-

¹⁸ Notice that, at the moment, the scrutable user model is read-only. Thus, subjects sent us the correct values by email.

¹⁹ Good et al. 1999 suggest that good MAE values should be near to 0.7, on a scale of values ranging from 0 to 5.

ing for events and related information. However they showed to be quite unaware about the content adaptation the system performs after some interactions. Thus we made the thumbs-up icons bigger, and put them before the stars. We also added a tooltip explaining the meaning of the icon (“this event is really suitable for you”). Finally, subjects were really attracted by the “similar users” function, and they liked to glance at other users’ public profiles. Users have also requested for more information about the others. They wanted to see what kinds of events their similar users and friends have inserted and bookmarked in order to find some implicit suggestions from people they trust in.

In order to further investigate the role of tags in the definition of the user model in iCITY, and their impact on the accuracy of recommendations, we carried out a second evaluation. We involved the same users as in the first evaluation and asked them to use iCITY for two weeks, every day. However, this time, they were explicitly asked not to use any tag. They could add events, click on events and rate them and add comments and information.

After the experimental period, they had to evaluate their scrutable user model and, if the scores did not match their real interest for the categories, they were asked to send the correct values back to the experimenters by email. All the users re-assigned their user model scores. To evaluate the accuracy, we calculated the MAE again. We obtained a medium MAE of 0.40 in a range from 0 to 1. We can easily observe that the new value is higher than the previous one (0.11). This demonstrates that, in iCITY, tags have an important role in the definition of the user model and in the accuracy of recommendations. In particular, they allow a faster updating of the user model, since they have a stronger impact with respect to other actions, and, most important, in a short period of time, they allow to define a user model that is more accurate than a user modeling strategy without tags. The same results would be reached in a longer period of time, in which users perform a greater number of actions on the system.

Some remarks: we have to underline that the promising results of this evaluation could suffer from some bias introduced by the way the experiments were carried out. First of all, subjects were chosen according to an *availability sampling approach* and this may cause a loss of external validity, namely the ability to make generalization about the obtained results. This could be solved by duplicating the experiments several times. Furthermore the subjects had to write an email in order to correct the wrong interest rating computed by the system. Some subjects could have not corrected the values since this would have caused an additional effort. Therefore, to extend the results, we are planning a larger evaluation with a greater number of real users

to collect further evidence on that. Moreover, not only for evaluation purposes, the scrutable user model will be modified to become more understandable and amendable for the user.

Finally, we examined all the tags in iCITY by analyzing the log files. We analyzed 437 tags used to annotate a total of 183 different events. We manually analyzed the meaning of the tags by classifying them into the categories described in Section 3. Considering the total number of tags, 321 (73%) were *free tags*, tags directly inserted by users, and 116 (27%) were *proposed tags*, tags proposed by the system. Therefore, this panel of users seems to prefer the use of personal tags rather than tags proposed by the system.

Considering the other categories, we found the following results:

- 212 (49%) out of 437 are *general tags*. Among them, general tags correspond not only to the upper classes of our ontology, but also to their subclasses. For instance, users group events by means of concepts such as *rock music*;
- 114 (26%) out of 437 are *specific tags*. Users often use very specific tags to describe the events, such as the singer’s name of a group in concert, the author’s name of a theater performance, etc. Other times they use the concepts an event title evokes (e.g. freedom, dreams, revenge, etc.). Users also tend to use the event title to tag the event;
- 58 (13%) out of 437 are *spatial tags*. Most of them describe the location where the event takes place i.e., the theater name, or the name of the street of the location;
- 28 (6%) out of 437 are *unknown tags*. The majority of them are unhyphenated compound words like “gem-boy”, or “noise-alternative”;
- 19 (4%) out of 437 are *subjective tags*. All of them are positive judgements, or adjective describing the event, i.e. “introspective”;
- 6 (1%) out of 437 are *synonym tags*.

This analysis led us to make some considerations about the user modeling approach adopted in iCITY. Users want to be free. Thus we should extend the lexicon of the controlled vocabulary, so that the system could analyze the meaning of inserted tags more specifically. Moreover we found that it is better not to classify tags corresponding to subclasses of our ontology as specific tags, since they are often used to group resources, not to describe them. As another improvement, most common event places and their address should be added to our

dictionary, so that the system could be able to recognize them. Finally the system should immediately recommend the event title, since it is often used to describe the event.

6. Related Works

Research topics related to our framework include recommender systems, mobile guides, user modeling systems, systems which use ontologies and Semantic Web techniques, Web 2.0-based systems which use social annotation and folksonomies, and systems in the field of cultural heritage. The first part of this section will briefly sketch the relation of our model with all these areas, taking them into account one by one. Then, the focus will be moved to works and projects which combine the above-mentioned areas, providing comparisons with our model. The last part of the section will deal with related works in the cultural heritage field.

PAPERRE is, first of all, a framework for recommender systems. Recommenders are often classified (Burke, 2002; Herlocker et al., 2004) into content-based recommenders (Aroyo et al., 2007; Cena et al., 2006), collaborative recommenders (Resnick et al., 1994; Shardanand and Maes, 1995) and hybrid recommenders (Balabanović and Shoham, 1997; Schwab and Koychev, 2000). PAPERRE and iCITY follow the content-based approach, since the recommendation is the result of the match between the attributes of the content and the attributes of the user profile. Recommender systems also share features with mobile guides (Stock et al., 2007; Cheverst et al., 2000; Rinner and Raubal, 2005; Pospischil et al., 2003) and in particular with tourism recommender mobile guides. GUIDE (Cheverst et al., 2000) is one of the first mobile guides whose features (e.g. adaptation to user location, user interests, time of the day) influenced the subsequent works in the field, and our framework too. Lol@ (Pospischil et al., 2003) is an example of a mobile guide that is adaptive toward the user device. Another related work is that of Rinner and Raubal (2005), which regards a location-based service that offers different criteria for calculating the score of items to recommend. With respect to these systems, the peculiarity of our approach mainly concerns the use of tagging as a feedback to improve the user model and thus to refine recommendations. As seen in Section 5, user models inferred using tags pick up better MAE values than user models calculated without taking them into account. To the best of our knowledge, up to now no other implemented content-based recommender exploits tagging to improve recommendations. Instead tags are used to im-

prove recommendations based on collaborative filtering and similarity measures (Zhao et al., 2008; Szomszor et al., 2007).

As regards the research area of user modeling, our framework follows the approach of explicit and scrutable representation of user modeling features. Related works are CHIP (Aroyo et al., 2007), which combines user models and scrutability and the works of Kay (2001, 2006), which first introduced the issue of scrutability. “The Adaptive Web”, a state of the art survey recently edited by Peter Brusilovsky, Alfred Kobsa and Wolfgang Nejdl (2007), contains comprehensive reviews about all the above-mentioned research areas.

Finally, considering the topic of knowledge representation, our framework is related to both current main approaches aimed at semantically enriching the knowledge of the system: the approach based on ontologies/taxonomies (Berners-Lee et al., 2006; Antoniou and van Harmelen, 2004) and the approach based on folksonomies (Golder and Huberman, 2006; Marlow et al., 2006; VanderWal, 2004; Macgregor and Mcculloch, 2006). In the research community, what is typical is the use of ontologies to represent knowledge (De Bra et al., 2004; Aroyo et al., 2007). An example is CHIME (De Bra et al., 2004), which represents user models by means of ontologies, and CHIP (Aroyo et al., 2007), which uses ontology-based user models. On the Web, examples of popular sites that use tagging and folksonomies are Delicious, Flickr, Eventful and CiteULike. Eventful is the closest to the iCITY prototype, since it provides a collection of events in the world and takes into account the user position.

The main contribution of our framework is the merge of all these topics. In the field of adaptive hypermedia, some of the works that merge adaptation and Web 2.0 techniques use social annotation to provide social navigation support, based on the direct and indirect behaviour of other people (Dourish and Chalmers, 1994; Svensson et al., 2001). Web 2.0 and social annotation naturally extend this practice, providing “a mix of both direct (intentional) navigational advice as well as indirect (inferred) advice based on the collective public behavior” (Millen and Feinberg, 2006). One of the most active groups that work on social adaptive annotation-based navigational support is the one led by Brusilovsky (Farzan and Brusilovsky, 2006a; Ahn et al., 2006). Particularly interesting is AnnotEd, an application to support browsing collections of linked educational resources on the Web. AnnotatEd adds adaptive visual cues to page links on the basis of traffic-based metrics and social annotations. Comparing this system with our model, we can observe that both of them provide navigation adaptation by means of item annotation (see Section 4.2). However, they are different from several points of view: open corpus resources (AnnotatEd) vs resources

restricted to information provided by the system and added by users (PAPERE/iCITY); user model composed of primary data, such as log activity and annotations, vs user model composed of primary data and inferred data; and, finally, adaptation based on link and item annotation vs adaptation based on recommendation of items, suggestion of similar users and item annotation as well.

A completely different approach combining Web 2.0 and personalization is related to the work of Bateman et al. (2006), which presents a framework inherent in our model with respect to the idea of linking tags to taxonomies. Bateman et al. propose a framework, CommonFolks, for integrating social tagging into a natural language ontology (WordNet), where students are led to annotate objects in a way that the added metadata can be machine processable. This approach is very interesting since it allows to overcome the problem of the lack of meaning in tags (Shirky, 2005), a problem particularly relevant in tagging-based recommenders. Considering our model, the main similarity is that in both cases domain ontology and tags are linked and this connection is achieved by supporting the tagging process²⁰. Szomszor et al. (2007) offer another example of integration of ontologies and folksonomies. The Movie Recommender they present provides predictions of rating for unseen movies. Both iCITY and the Movie Recommender define a domain ontology and store the tags inserted by users. However, there are relevant differences to point out. In our prototype, differently from the Movie Recommender, tags are associated to the resource but also to the user who inserted them, in order to perform user modeling inferences. Moreover the Movie Recommender uses a collaborative approach based on tag similarities and is not interested in the meaning of tags.

Considering the use of tagging for *user modeling*, we have found that, up to now, very few systems have explored this kind of feedback to extend the model of the user. Van Setten et al. (2006) suggest that annotations can become part of the user profile as “an indication of his perspective on the content collection and interest in the annotated object”. Comparing this approach to our model, we can observe that it does not deal with the issue of how to exploit tagging in order to infer knowledge about the user, which is the core of our model. They simply suggest that this profile can be used for recommendation, making use of techniques such as collaborative filtering or case-based reasoning. Nauerz et al. (2007) propose a similar approach sketching some similarity metrics aimed at modeling users and providing recommendations.

²⁰ As explained in Section 4.5, in iCITY it is achieved by suggesting tags which belong to vocabulary linked to the domain taxonomy.

Finally, let us analyze works related to the topic of cultural heritage. Several works are related to our framework and to iCITY with respect to the adoption of ontologies, folksonomies and adaptivity to face the issues we enlightened in Section 1, summarizable as improving the accessibility, interest and elicitation of cultural heritage resources. Many authors associate the problem of accessibility to the modality of classification of resources. For example, the works of Smith (2006), Trant (2006), Chun et al. (2006) analyze the field of museums, explaining that art collections typically suffer from the problem that non-specialists have a hard time retrieving resources classified by specialists. They claim that tagging and folksonomies can offer a way to overcome the well-known *Vocabulary Problem*, formulated by Furnas et al. (1987). In particular, the results of the experiments in the Steve.museum project²¹ (Bearman and Trant, 2005; Trant, 2006; Chun and Wyman, 2006), show that when users add tags about artworks, they insert information about the subject of the artwork (such as, in case of painting, the people and place represented, the suggested ideas, emotions, etc.), while experts provide only “external” information regarding the authors, the historical period, materials and so on. Moreover it emerged that users are generally keen to leave a trace of what they think and feel. Thus, the addition of metadata containing the user’s point of view can be a pleasant experience for “author-users” and a useful opportunity for other users to find resources they are interested in. The users’ point of view can be also useful to index cultural heritage resources with localistic labels related, for instance, to historical local events, and to address the issue that the meaning of terms people use may change over the time. Several works (Abbott, 2004; Chan, 2007; Szomszor et al., 2007) suggest that the best solution for resources accessibility would be to integrate the subjective perspective of users within traditional classification systems. In this way the benefits of both the approaches could be exploited, limiting their respective problems (this is the approach followed in our framework). A paper which is often cited is the one by Golder and Huberman (2006). The authors synthetically describe the features of these two approaches, while other works deal with the same subject more extensively (Mathes, 2004; Shirky, 2005; VanderWal, 2004).

Another relevant issue related to the cultural heritage domain concerns the use of strategies to make people aware that they are interested

²¹ The Steve.Museum Social Tagging Project (<http://www.steve.museum/>) is a collaboration among a set of important art museums (such as the Metropolitan Museum of Art and the Guggenheim Museum) to explore the advantages coming from the use of social tagging, to study its effectiveness in improving the access to museum collections and to understand what motivates individuals to tag.

in specific resources, avoiding cursory glances at resources. While the previous issue was related to resource accessibility and retrieval, the current one is related to resource filtering and recommendation. CHIP (Aroyo et al., 2007; Cramer et al., 2008) is a project in the domain of museums where authors move on from this problem and try to apply techniques from semantic annotation and user profiling to make the visits of the museum and of the museum's web site more selective. Thus, it is relevant with respect to the personalization in real exploration of cultural heritage. In the system, people are asked to rate artwork items. Then the system looks for other artworks and related topics in the semantic network in order to recommend items and suggest guided tours. They are also working on introducing tools to allow for the active engagement of users with artworks in the collection (Aroyo et al., 2007). The system is also important since it offers the users insight in the reasons why a particular recommendation has been made. Experimental results showed that such transparency of the reasons for recommendations increased acceptance of recommendations themselves (Cramer et al., 2008).

Concerning the personalization in real exploration of cultural heritage, we also mention Carletto (Damiano et al., 2008), a character-based adaptive guide for mobile devices. Carletto has been developed for an historical site and used by real users during the three days opening of that site. The distinctive features of the application are the drama-based approach in the information presentation, and the lightweight adaptation strategy. Concerning in particular adaptation features, Carletto is adaptive with respect to the location, relying on a wireless network infrastructure in the historical site. Moreover the selection of contents keeps track of the interaction history, in order to deliver the appropriate quantity and type of informative items, and to manage the given/new distinction in discourse.

A further example of personalization in real exploration of cultural heritage is the LISTEN project (Zimmermann and Lorenz, 2008). LISTEN makes use of inherent "everyday" integration of the aural and visual perception, developing a tailored, immersive audio augmented environment for the visitors of art exhibitions. LISTEN provides the user with an audio presentation which is adapted to the users contexts. The auditory rendering process takes into account the users current position, the head orientation, her interests, preferences, and motion in order to seamlessly integrate the virtual scene with the real one. Speech, music and sound effects are dynamically arranged to form an individualized and situated soundscape offering information related to visual objects as well as creating context-specific atmospheres.

Exploiting the emulative behaviour of users is another way to increase the interest in cultural heritage domains. Users that have no specific knowledge or interest in a domain can be intrigued by items that many users view or tag (see the success of sites like You Tube) and a consequence is the increase of the possibility of serendipity. Chan (2007) deals about this while presenting the case of the Powerhouse Museum in Sydney. He defines serendipity as the art of chancing upon things and making unexpected discoveries, allowing users to come across items that are potentially interested in, but that they were unaware of prior to visiting.

The last related issue we would like to discuss regards the detection, elicitation and consequently enhancement of cultural heritage resources. The cultural heritage is often scattered and spread over the territory, without a proper classification and publicity. Consider, for example, distributed collections, mural works, buildings of multi-cultural interest not directly classifiable into typical categories, street performances, etc. To detect such elements of heritage, the basic principle of our framework is to involve users, inhabitants, citizens, etc. and offer them tools to elicit this knowledge easily. Related to this issue, “Share Your Heritage”²² is an interesting project that uses stories told by users to link people to sites with the outcomes of making people aware of the wealth of cultural heritage resources in the region and enhancing the visitor experience. Gorla et al. (2003) talk about the “narrative dimension” of a place, relating it to the possibility of enriching a cultural heritage resource with some episodes, stories and anecdotes related to its history.

In the last years, several projects related to this issue were presented at the main conferences²³ on Cultural Heritage and Museums on the Web. Fisher et al. (2008) deal with the possibility to enrich artwork catalogues through visitor narratives. Mannion (2008) suggests that, in addition to provide a supplementary vocabulary to describe and access works of art, as stated by the already-mentioned work of Chun et al. (2006), tags can convey the different perception of works from viewers of diverse cultures (in their study they collected tags on Tibetan artworks from Tibetans and Westerners in Switzerland and New York). Oates (2008) highlights that social annotations often produce cross-references of artworks with other information on the Internet, as it emerged by analyzing the annotations provided by Flickr users to images from the Library of Congress Prints and Photographs. Chan (2007) focuses his attention on the opportunity to use social annotation to enrich the

²² <http://www.ruralaction.org/>

²³ <http://www.archimuse.com/>

information regarding heritage resources already available in official circuits but poor in description and details.

Supporting the user by supplying an environment for *peer production*²⁴ (Benkler, 2006), as in our framework, increases the possibility to facing these issues. Other examples of frameworks that use Web 2.0 tools for peer production in the cultural heritage domain are iAKS (Watrall and Siarto, 2008), a knowledge management system to enrich and manage archaeological resources, and Open Context (Kansa and Kansa, 2007), a Web 2.0-based tool to manage archaeological and museum collection datasets. In this case, the community of authors is composed by researchers and collections managers. Finally, related to the use of bookmarking and annotation, we can mention the use of these tools in several museums to allow visitors to keep track of interesting information for later use, accessed on-line in a personal page or via e-mail and shared with others or not (e.g., Beler et al. 2004). A recent study (Filippini-Fantoni and Bowen, 2007) analyzes a number of museums which follow this approach to evaluate benefits and weaknesses. Results show that bookmarking is used just by a limited subset of users. However this small but important core of visitors find such facilities extremely worthwhile. This study can be an important reference to consider when providing similar kinds of tools, as in our framework. An extensive review of projects which use Web 2.0 tools, like tagging, blogs, social bookmarking, etc., in the cultural heritage domain, is provided in Middleton and Lee (2007).

7. Discussion and conclusions

In this paper we have presented PAPERRE, a framework for the integration of the Web 2.0 paradigm, especially social annotation, with user modeling, showing that both the methodologies can benefit from a proper integration, and explaining why the combination of these approaches within cultural heritage applications is particularly significant for the promotion of cultural heritage itself. We also presented iCITY, a prototype implementation of the PAPERRE model for accessing cultural heritage resources about the city of Turin. The evaluation of iCITY, reported in the paper, showed that users seem to appreciate the system and especially its social and community features.

The main contribution of the framework is the proposal of a new approach to user modeling and to content-based recommender systems.

²⁴ Peer production is a term coined by Yochai Benkler (2006) to describe a new economic model characterized by the voluntary and community-regulated production of an intellectual work and the term is frequently applied to the Web 2.0 paradigm.

The innovation consists in exploiting the natural user activity of tagging to build and/or improve the user model, and thus improve recommendations. Moreover, other contributions are the use of tags (and folksonomies derived from tags) to create personalized ways of navigating contents in a bottom-up approach and the use of user modeling to support users in tagging, in creating contents and in navigation.

However, our approach has some limitations. Thus, we wish to conclude the paper discussing the bounds of the framework and its conditions of applicability. First of all, considering the projects mentioned in Section 6, it emerges that all the tag-based examples regard applications for cataloguing, promoting and recommending cultural heritage resources. Indeed, tagging is an approach which offers personal ways of classifying and recovering resources and, thus, it is more useful in the stage of retrieval and recommendation, than in the stage of watching and enjoyment of the resources. Since our framework introduces a tag-based approach into a content-based recommender, it intrinsically inherits this bound.

Another issue regards the possibility of applying the framework in conditions of one-time use, as often happens in the tourist domain, or whether it works better in conditions of frequent use. The employment of long-term user models and the emphasis on the community of users which share similar interests lead us to state that the latter condition allows to take advantage of all the solutions provided by the framework, even if it can be used also in conditions of one-time use. iCITY, for instance, is mainly conceived for people living a territory who create a community of users. However, iCITY can be used by one-time tourist as well, and tourists can provide an interesting contribution to enriching system knowledge, since they view from outside the cultural heritage of the territory. One of the central issues for the framework is how to induce people to participate in the content production. Does the well-known paradox of the active user (Carroll and Rosson, 1987) affect annotations in the same way as it regards the provision of explicit feedback? Farzan and Brusilovsky (2006b), following some ideas of Bretzke and Vassileva (2003), state that a reward mechanism is necessary in order to obtain user feedback implicitly, as part of the users natural interaction with the system. The success of tagging web sites like del.icio.us (where tagging allows users to organize their bookmarks) or flickr.com (where tagging allows to share photos) confirms it. However, there are lots of examples where people provide comments and annotations without a direct and tangible reward mechanism. The boom of *opinion websites*, like, for example, dooyoo.com and epinions.com, where users provide comments just for the wish of sharing their experiences and being useful to others, show

that the dynamics for the provision of annotations are complex and rich. Several works investigate the reasons for passive attendance, or lurking behavior, vs active participation in online communities (Preece et al., 2004; Soroka and Rafaeli, 2006; Soroka and Jacovi, 2004). Several answers were given - the need to learn before participating, the use of the common goods without contributing to it, the fear to be mocked, etc. However, the most relevant conclusion of these works is that all users, opportunely stimulated, can de-lurk. In particular, some studies by Soroka and Rafaeli (2006) found that the probability a lurker becomes an active participant is correlated to the time she invests in reading over a short period of time. They suggest that designers should make the first experiences in the community as pleasant, unobtrusive and immediate as possible. Moreover, it has been found that a critical mass of quality content rather than of users can also contribute to the successful spread of a virtual community (Soroka and Jacovi, 2004). These results are relevant for our discussion since they provide a confirmation of the validity of several choices we made in our framework and in iCITY: *i*) the provision of quality content by the system (e.g. RSS feed from the city of Torino's cultural portal) and the filtering of bad contents; *ii*) the attention toward lurkers, by defining user modeling algorithms that exploit different kinds of implicit feedbacks and not just tagging; *iii*) the accurate design of the user interface and of the different modalities to navigate items, aimed at making the experience pleasant and the functionality learnable.

Some final considerations regard the iCITY prototype specifically. First of all, there is a need for further investigation regarding social dynamics - e.g. the definition and evolution of *information credibility* with reference to user-generated contents and network dynamics. Among other subjects, it could be useful to determine if openness to different information sources, including potentially critical contributions, enhances the perception that the application provides accurate, credible, trustworthy and unbiased information.

Another aspect that has to be considered concerns comments. Comments are "free", in the sense that they can be written as free texts and they can provide both a piece of information about an event or a user opinion. At the current stage, our prototype is not equipped with Natural Language Processing capabilities, and thus these kinds of annotations are not exploited to learn user interest or attitudes. Obviously, this is a limitation, since comments can be an important source of information about the user. A possible extension of the iCITY system would consist in enabling the Content Insertion Manager to process free text, in order to exploit free comments (e.g., user opinions), even in the case of misspelled words or regional variants. In the same

direction, the evaluation of the iCITY prototype highlighted that users often tag by using personal labels. This suggests that the inclusion of an extended lexicon, or of a Natural Language Processing module that allows the system to understand the meaning of personal labels, could be useful.

8. Acknowledgments

We would like to thank Franco Carcillo, the Municipality of Torino and CSP for the support in the project. The research has also been partially supported by the Italian Ministry of Education, University and Research (MIUR) under the PRIN 2005 project “Adaptive, Context-Aware, Multimedia Guides on Mobile Devices”.

References

- Abbott, R.: 2004, ‘Subjectivity as a concern for information science: a Popperian perspective’. *Journal of Information Science* **30**(1), 95–106.
- Abowd, G. and E. Mynatt: 2000, ‘Charting past, present and future research in ubiquitous computing’. *ACM Transactions on Computer-Human Interaction, Special Issue on HCI in the new Millennium* **7**(1), 29–58.
- Ahn, J., R. Farzan, and P. Brusilovsky: 2006, ‘A Two-Level Adaptive Visualization for Information Access to Open-Corpus Educational Resources’. In: *Proceedings of the Workshop on the Social Navigation and Community based Adaptation Technologies at AH 2006*. Dublin, Ireland, pp. 497–505.
- Antoniou, G. and F. van Harmelen: 2004, *A Semantic Web Primer*. Cambridge, MA: The MIT Press.
- Aroyo, L., R. Brussee, P. Gorgels, L. Rutledge, N. Stash, and Y. Wang: 2007, ‘Personalized Museum Experience: The Rijksmuseum Use Case’. In: J. Trant and D. Bearman (eds.): *Proceedings of Museums on the Web 2007: Selected Papers from an international conference*. San Francisco, CA, USA. <http://www.archimuse.com/mw2007/papers/aroyo/aroyo.html>, accessed May, 2008.
- Balabanović, M. and Y. Shoham: 1997, ‘Fab: content-based, collaborative recommendation’. *Communications of the ACM* **40**(3), 66–72.
- Bateman, S., C. Brooks, and G. McCalla: 2006, ‘Collaborative Tagging Approaches for Onto-logical Metadata in Adaptive E-Learning Systems’. In: *Proceedings of the Workshop on Applications of Semantic Web Technologies for E-Learning at AH 2006*. Dublin, Ireland, pp. 3–12.
- Bearman, D. and J. Trant: 2005, ‘Social Terminology Enhancement through Vernacular Engagement: Exploring Collaborative Annotation to Encourage Interaction with Museum Collections’. *D-Lib Magazine* **11**(9). <http://www.dlib.org/dlib/september05/bearman/09bearman.html>, accessed May, 2008.
- Beler, A., A. Borda, J. P. Bowen, and S. Filippini-Fantoni: 2004, ‘The Building of Online Communities: An approach for learning organizations, with

- a particular focus on the museum sector'. In: *Proceedings of EVA 2004 London Conference, University College London, The Institute of Archaeology*. <http://arxiv.org/ftp/cs/papers/0409/0409055.pdf>, accessed May 2008.
- Benkler, Y.: 2006, *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. New Haven, Conn: Yale University Press.
- Berners-Lee, T., W. Hall, J. A. Hendler, K. O'Hara, N. Shadbolt, and D. J. Weitzner: 2006, *A Framework for Web Science. Foundations and Trends in Web Science*. Hanover, MA: Now Publishers Inc.
- Bretzke, H. and J. Vassileva: 2003, 'Motivating Cooperation on Peer to Peer Networks'. In: P. Brusilovsky, A. T. Corbett, and F. de Rosis (eds.): *Proceedings of User Modeling 2003*, Vol. 2702 of *Lecture Notes in Computer Science*. Pittsburgh, Johnstown, USA, pp. 218–227.
- Brusilovsky, P.: 1996, 'Methods and Techniques of Adaptive Hypermedia.'. *User Modeling and User-Adapted Interaction* **6**(2-3), 87–129.
- Brusilovsky, P., A. Kobsa, and W. Nejdl (eds.): 2007, *The Adaptive Web, Methods and Strategies of Web Personalization*, Vol. 4321 of *Lecture Notes in Computer Science*. Berlin: Springer.
- Brusilovsky, P. and M. T. Maybury: 2002, 'From adaptive hypermedia to the adaptive web'. *Communications of the ACM, Special Issue: The Adaptive Web* **45**(5), 30–33.
- Burke, R.: 2002, 'Hybrid recommender systems: Survey and experiments'. *User Modeling and User-Adapted Interaction* **12**(4), 331–370.
- Carmagnola, F., F. Cena, O. Cortassa, C. Gena, and I. Torre: 2007, 'Towards a Tag-Based User Model: How Can User Model Benefit from Tags?'. In: C. Conati, K. F. McCoy, and G. Paliouras (eds.): *Proceedings of User Modeling 2007*, Vol. 4511 of *Lecture Notes in Computer Science*. Corfu, Greece, pp. 445–449.
- Carroll, J. M. and M. B. Rosson: 1987, 'Paradox of the active user'. In: *Interfacing thought: cognitive aspects of human-computer interaction*. Cambridge, MA: The MIT Press, pp. 80–111.
- Cena, F., L. Console, C. Gena, A. Goy, G. Levi, S. Modeo, and I. Torre: 2006, 'Integrating heterogeneous adaptation techniques to build a flexible and usable mobile tourist guide'. *AI Communications* **19**(4), 369–384.
- Cena, F., C. Gena, and S. Modeo: 2005, 'How to Communicate Recommendations? Evaluation of an Adaptive Annotation Technique'. In: M. F. Costabile and F. Paternò (eds.): *Proceedings of INTERACT 2005*, Vol. 3585 of *Lecture Notes in Computer Science*. Rome, Italy, pp. 1030–1033.
- Chan, S.: 2007, 'Tagging and Searching? Serendipity and museum collection databases'. In: J. Trant and D. Bearman (eds.): *Proceedings of Museums on the Web 2007: Selected Papers from an international conference*. San Francisco, CA, USA. <http://www.archimuse.com/mw2007/papers/chan/chan.html>, accessed May, 2008.
- Cheverst, K., N. Davies, K. Mitchell, A. Friday, and C. Efstratiou: 2000, 'Developing a context-aware electronic tourist guide: some issues and experiences'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2000*. The Hague, The Netherlands, pp. 17–24.
- Chun, S. Hiwiler, D. T. J. and B. Wyman: 2006, 'Steve.museum: An On-going Experiment in Social Tagging, Folksonomy, and Museums'. In: J. Trant and D. Bearman (eds.): *Proceedings of Museums on the Web 2006: Selected Papers from an international conference*. Albuquerque, New Mexico, USA. <http://www.archimuse.com/mw2006/papers/wyman/wyman.html>, accessed May, 2008.

- Cramer, H., V. Evers, S. Ramlal, M. van Someren, B. Wielinga, L. Rutledge, N. Stash, and L. Aroyo: 2008, 'The effects of transparency on trust in and acceptance of a content-based art recommender'. *This issue*.
- Damiano, R., C. Gena, V. Lombardo, F. Nunnari, and A. Pizzo: 2008, 'A stroll with Carletto. Adaptation in drama-based, interactive tours with virtual characters'. *This issue*.
- De Bra, P., L. Aroyo, and V. Chepegin: 2004, 'The Next Big Thing: Adaptive Web-Based Systems'. *Journal of Digital Information* **5**(247). <http://journals.tdl.org/jodi/article/view/jodi-135/122>, accessed March, 2008.
- Dey, A. and D. Abowd: 2000, 'Towards a better understanding of context and context-awareness'. In: *Proceedings of the Workshop the What, Who, Where, When and How of Context - Awareness at CHI 2000*. The Hague, The Netherlands.
- Dourish, P. and M. Chalmers: 1994, 'Running out of Space: Models of Information Navigation.'. In: G. Cockton, S. W. Draper, and G. R. S. Weir (eds.): *People and Computers IX, Proceedings of HCI '94*. Glasgow, Scotland.
- Farzan, R. and P. Brusilovsky: 2006a, 'AnnotatEd: A Social Navigation and Annotation Service for Web-based Educational Resources'. In: *Proceedings of the World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education, E-Learn 2006*. Honolulu, Hawaii, pp. 2794–2802.
- Farzan, R. and P. Brusilovsky: 2006b, 'Social Navigation Support in a Course Recommendation System'. in (Wade et al., 2006), pp. 91–100.
- Fellbaum, C.: 1998, *WordNet: An Electronic Lexical Database*. Cambridge, MA: The MIT Press.
- Filippini-Fantoni, S. and J. Bowen: 2007, 'Bookmarking In Museums: Extending The Museum Experience Beyond the Visit?'. In: *Proceedings of the Int. Conf. on Museums and the Web 2007*. <http://www.archimuse.com/mw2007/papers/filippini-fantoni/filippini-fantoni.html>, accessed May 2008.
- Fisher, M., B. Twiss-Garrity, and A. Sastre: 2008, 'The Art of Storytelling: Enriching Art Museum Exhibits and Education through Visitor Narratives'. In: *Proceedings of the Int. Conf. on Museums and the Web 2008*. <http://www.archimuse.com/mw2008/papers/fisher/fisher.html>, accessed May 2008.
- Furnas, G. W., T. K. Landauer, L. M. Gomez, and S. T. Dumais: 1987, 'The vocabulary problem in human-system communication'. *Communications of the ACM* **30**(11), 964–971.
- Gena, C. and S. Weibelzahl: 2007, 'Usability Engineering for the Adaptive Web'. in (Brusilovsky et al., 2007), pp. 720–762.
- Golder, S. and B. A. Huberman: 2006, 'Usage Patterns of Collaborative Tagging Systems'. *Journal of Information Science* **32**(2), 198–208.
- Good, N., J. B. Schafer, J. A. Konstan, A. Borchers, B. M. Sarwar, J. L. Herlocker, and J. Riedl: 1999, 'Combining Collaborative Filtering with Personal Agents for Better Recommendations'. In: *Proceedings of the Sixteenth National Conference on Artificial Intelligence and Eleventh Conference on Innovative Applications of Artificial Intelligence, AAAI/IAAI 1999*. Orlando, Florida, pp. 439–446.
- Gorla, S., A. Ciocca, A. Cioffi, and E. Delle Donne: 2003, 'Manuale di didattica museale on-line'. *2F Multimedia - MULTIMEDIA EDUCATION, Firenze*.
- Herlocker, J. L., J. A. Konstan, L. G. Terveen, and J. T. Riedl: 2004, 'Evaluating collaborative filtering recommender systems'. *ACM Transactions on Information Systems* **22**(1), 5–53.

- Jameson, A.: 2006, 'Adaptive Interfaces and Agents'. In: J. A. Jacko and A. Sears (eds.): *Human-computer interaction handbook (2nd ed.)*. Mahwah, NJ: Erlbaum, pp. 305–330.
- Kansa, E. and S. W. Kansa: 2007, 'Open Context: Collaborative Data Publication to Bridge Field Research and Museum Collections'. In: *Proceedings of the International Cultural Heritage Informatics Meeting (ICHIM07)*. <http://www.archimuse.com/ichim07/papers/kansa/kansa.html>, accessed May 2008.
- Kay, J.: 2006, 'Scrutable Adaptation: Because We Can and Must'. in (Wade et al., 2006), pp. 11–19.
- Kay, J., R. J. Kummerfeld, and P. Lauder: 2001, 'Foundations for personalised documents: a scrutable user Model Server'. In: A. M. Vercoustre and J. Kay (eds.): *Proceedings of the Australian Document Computing Symposium, ADCS 2001*. Coff's Harbour, Australia, pp. 43–50.
- Kobsa, A., J. Koenemann, and W. Pohl: 2001, 'Personalized Hypermedia Presentation Techniques for Improving Online Customer Relationships.'. *The Knowledge Engineering Review* **16**(2), 111–155.
- Macgregor, G. and E. McCulloch: 2006, 'Collaborative Tagging as a Knowledge Organisation and Resource Discovery Tool'. *Library Review* **55**(5), 291–300.
- Mannion, S.: 2008, 'Seeing Tibetan Art Through Social Tags'. In: *Proceedings of the Int. Conf. on Museums and the Web 2008*. <http://www.archimuse.com/mw2008/papers/mannion/mannion.html>, accessed May, 2008.
- Marlow, C., M. Naaman, D. Boyd, and M. Davis: 2006, 'HT06, tagging paper, taxonomy, Flickr, academic article, to read'. In: *Proceedings of the Seventeenth Conference on Hypertext and Hypermedia, HYPERTEXT 2006*. Odense, Denmark, pp. 31–40.
- Mathes, A.: 2004, 'Folksonomies - Cooperative Classification and Communication Through Shared Metadata'. <http://www.adammathes.com/academic/computer-mediated-communication/folksonomies.html>, accessed January 2008.
- Middleton, M. R. and J. M. Lee: 2007, 'Cultural institutions and Web 2.0'. In: *Proceedings of the Fourth Seminar on Research Applications in Information and Library Studies (RAILS 4)*. <http://eprints.qut.edu.au/archive/00010808/01/Cultural.Institutions.and.Web.2.0.pdf>, accessed May 2008.
- Millen, D. R. and J. Feinberg: 2006, 'Using Social Tagging to Improve Social Navigation'. In: *Proceedings of the Workshop on the Social Navigation and Community-Based Adaptation Technologies at AH 2006*. Dublin, Ireland, pp. 532–541.
- Nauerz, A., S. Pietschmann, and R. Pietzsch: 2007, 'Collaborative annotation-driven adaptation in web portals'. In: *Proceedings of the Eighteen Conference on Hypertext and Hypermedia, HYPERTEXT 2007*. Manchester, UK, pp. 155–156.
- Nielsen, J.: 1999, *Designing Web Usability: The Practice of Simplicity*. Harlow, Essex: New Riders Press.
- Oates, G.: 2008, 'The Commons on Flickr: A Primer'. In: *Proceedings of the Int. Conf. on Museums and the Web 2008*. <http://www.archimuse.com/mw2008/papers/oates/oates.html>, accessed May 2008.
- O'Reilly, T.: 2005, 'What Is Web 2.0. Design Patterns and Business Models for the Next Generation of Software'.

- <http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>, accessed January 2008.
- Pospischil, G., M. Umlauft, and E. Michlmayr: 2003, 'Designing LoL@, a Mobile Tourist Guide for UMTS'. *Information Technology and Tourism* **5**(3), 151–164.
- Preece, J., B. Nonnecke, and D. Andrews: 2004, 'The top five reasons for lurking: improving community experiences for everyone.'. *Computers in Human Behavior* **20**(2), 201–223.
- Resnick, P., N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl: 1994, 'GroupLens: An open architecture for collaborative filtering of netnews'. In: *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW 1994*. Chapel Hill, NC, USA, pp. 175–186.
- Rinner, C. and M. Raubal: 2005, 'Personalized Multi-Criteria Decision Strategies in Location-Based Decision Support'. *Journal of Geographic Information Sciences* **11**(1), 61–68.
- Royce, S. A. and B. C. Straits: 1999, *Approaches to Social Research (3rd Edition)*. New York: Oxford University Press.
- Schafer, B., D. Frankowski, J. Herlocker, and S. Sen: 2007, 'Collaborative Filtering Recommender Systems'. in (Brusilovsky et al., 2007).
- Schwab, I., P. and W. Koychev, I.: 2000, 'Learning to recommend from positive evidence'. In: *Proceedings of the International Conference on Intelligent User Interfaces, IUI 2000*. New Orleans, Louisiana, USA, pp. 241–247.
- Sen, S., S. K. Lam, A. M. Rashid, D. Cosley, D. Frankowski, J. Osterhouse, F. M. Harper, and J. Riedl: 2006, 'Tagging, communities, vocabulary, evolution'. In: *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW 2006*. Banff, Alberta, Canada, pp. 181–190.
- Shardanand, U. and P. Maes: 1995, 'Social information filtering: algorithms for automating word of mouth'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 1995*. Denver, Colorado, United States, pp. 210–217.
- Shirky, C.: 2005, 'Ontology is overrated: Categories, links, and tags'. http://www.shirky.com/writings/ontology_overrated.html, accessed January 2008.
- Smith, M. K.: 2006, 'Viewer Tagging In Art Museums: Comparisons To Concepts And Vocabularies Of Art Museum Visitors'. In: *Proceedings of the Seventeenth Workshop in Classification Research*. Austin, Texas, USA.
- Soroka, V. and M. Jacovi: 2004, 'The diffusion of reachOut: analysis and framework for the successful diffusion of collaboration technologies'. In: *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW 2004*. Chicago, Illinois, USA, pp. 314–323.
- Soroka, V. and S. Rafaeli: 2006, 'Invisible participants: how cultural capital relates to lurking behavior'. In: *Proceedings of the Fifteenth International Conference on World Wide Web, WWW 2006*. Edinburgh, Scotland, pp. 163–172.
- Stock, O., M. Zancanaro, P. Busetta, C. Callaway, A. Krüger, M. Kruppa, T. Kuflik, E. Not, and C. Rocchi: 2007, 'Adaptive, intelligent presentation of information for the museum visitor in PEACH'. *User Modeling and User-Adapted Interaction* **17**(3), 257–304.
- Strauss, A. and J. Corbin: 1990, *Basics of qualitative research: grounded theory procedures and techniques*. Newbury Park, CA: Sage Publications.
- Svensson, M., K. Höök, J. Laaksolahti, and A. Waern: 2001, 'Social navigation of food recipes'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2001*. Seattle, Washington, USA, pp. 341–348.

- Szomszor, M., C. Cattuto, H. Alani, K. Oâ€™Hara, A. Baldassarri, V. Loreto, and V. D. Servidio: 2007, ‘Folksonomies, the Semantic Web, and Movie Recommendation’. In: *Proceedings of the Fourth European Semantic Web Conference*. Innsbruck, Austria, pp. 71–85.
- Trant, J.: 2006, ‘Exploring the potential for social tagging and folksonomy in art museums: Proof of concept’. *New Review of Hypermedia and Multimedia* **12**(1), 83–105.
- van Setten, M., R. Brussee, H. van Vliet, L. Gazendam, Y. van Houten, and M. Veenstra: 2006, ‘On the Importance of “Who Tagged What”’. In: *Proceedings of the Workshop on the Social Navigation and Community based Adaptation Technologies at AH 2006*. Dublin, Ireland, pp. 552–561.
- VanderWal, T.: 2004, ‘Folksonomy Coinage and Definition’. <http://vanderwal.net/folksonomy.html>, accessed January 2008.
- Wade, V., H. Ashman, and B. Smyth (eds.): 2006, ‘Adaptive Hypermedia and Adaptive Web-Based Systems. Fourth International Conference AH 2006’, Vol. 4018 of *Lecture Notes in Computer Science*. Berlin:, Springer.
- Watrall, E. and J. Siarto: 2008, ‘iAKS: A Proposal for a Web 2.0 Archaeological Knowledge Management System’. In: *Proceedings of the International Cultural Heritage Informatics Meeting (ICHIM07)*. <http://www.archimuse.com/ichim07/papers/watrall/watrall.html>, accessed May 2008.
- Xu, Z., Y. Fu, J. Mao, and D. Su: 2006, ‘Towards the Semantic Web: Collaborative Tag Suggestions’. In: *Proceedings of the Collaborative Web Tagging Workshop at WWW 2006*. Edinburgh, Scotland. <http://www.semanticmetadata.net/hosted/taggingws-www2006-files/13.pdf>, accessed May 2008.
- Zhao, S., N. Du, A. Nauerz, X. Zhang, Q. Yuan, and R. Fu: 2008, ‘Improved Recommendation based on Collaborative Tagging Behaviors’. In: *Proceedings of the International Conference on Intelligent User Interfaces, IUI 2008*. Canary Islands, Spain, pp. 413–416.
- Zimmermann, A. and A. Lorenz: 2008, ‘LISTEN: A User-Adaptive Audio-Augmented Museum Guide’. *This issue*.

Authors’ vitae

Francesca Carmagnola is a short time researcher in Computer Science at the University of Torino. She received her PhD in 2007 in Communication Science (University of Torino), in the area of distributed and interoperable user models. Her current research activities address the user modelling, personalization of adaptive hypermedia systems, personalization of ubiquitous systems, infrastructures for adaptive web-based systems, personalized digital television, semantic knowledge representation and reasoning, semantic web.

Federica Cena is a short time researcher at the Department of Computer Science of the University of Torino. She received her PhD in 2007 in Communication Science (University of Torino) in 2007, in the area of

semantic web for distributed web applications. She is working on user modelling, personalization, ubiquitous computing applications and web services.

Luca Console (PhD 1991) is Full Professor at the Department of Computer Science of the University of Torino, working in the area of Knowledge Representation and Reasoning. His research focus is on Model-Based Reasoning (Model-based diagnosis) and on adaptive systems. He coordinated the iCITY project and contributed on aspects of knowledge representation and reasoning. He coordinated several national and international research projects and published more than 100 paper on International Journals, collections and proceedings of international conferences.

Omar Cortassa is a temporary researcher at Telecom Italia Lab, the Research and Development department of Telecom Italia. He got his master's degree in Computer Science in December 2007, University of Torino. Previously he worked as a researcher at the Department of Computer Science of the University of Torino. From March 2006 he is a member of iCITY Program where he works on system architecture definition and maintenance. His main interests address multi-tier architecture for web-based systems, service-oriented architecture, knowledge representation for adaptive systems, agent-based technologies, personalization and reasoning in the Semantic Web.

Cristina Gena is Researcher at the Department of Computer Science of the University of Torino, working in the area of intelligent user interfaces, and Assistant Professor at the School of Multimedia and Arts, University of Torino. She completed her Ph.D. in Communication Science (University of Torino) in 2003, with a thesis on the evaluation of user-adaptive systems. Her current research activities address user modeling, adaptive web systems and their evaluation, context-aware systems, semantic web, web 2.0, usability and interaction design.

Anna Goy is Researcher at the Department of Computer Science of the University of Torino, working in the area of web-based systems. She completed her Ph.D. in Cognitive Science (University of Torino) in 1998, in the area of lexical semantics. Since many years she works on distributed web-based applications, including adaptive hypermedia, context-aware systems, and Web technologies.

Ilaria Torre is Researcher at the Department of Computer Science of the University of Torino. She works in the area of intelligent user

interfaces and is Assistant Professor at the School of Communication Sciences, University of Torino. She completed her Ph.D. in Communication Science and Project (University of Torino) in 2003, with a thesis on the process of user modeling in adaptive hypermedia systems. Her current research activities address user modeling, context-aware adaptive systems, knowledge representation, semantic web, web 2.0.

Andrea Toso is Web area supervisor at Wedoo, a web agency in Torino, working in the area of social media applications. At the time of iCITY development, he was a Project Manager for CSP, a research centre in Torino. Since many years he works on online web communities environment, including context-aware systems, social media, user generated geolocalized content.

Fabiana Venero is, at present, a PhD Student at the Department of Computer Science of the University of Torino. After graduating in 2005, she worked for two years as a Web analyst and Human-Machine Interaction expert for CSP, a research centre in Torino. Her main interests lie in the areas of Human-Machine Interaction and Interaction Design.