A SEMANTICALLY-ENHANCED PERSONALIZATION FRAMEWORK FOR KNOWLEDGE-DRIVEN MEDIA SERVICES

David Vallet[1], Phivos Mylonas[2], Miguel A. Corella[1], José M. Fuentes[1],
Pablo Castells[1], and Yannis Avrithis[2]

[1]*Universidad Autónoma de Madrid, Escuela Politécnica Superior*
Campus de Cantoblanco, 28049 Madrid, Spain
{david.vallet, miguel.corella, chema.fuentes, pablo.castells}@uam.es

[2]*National Technical University of Athens, School of Electrical and Computer Engineering*
GR-15773 Zographou, Athens, Greece
{fmylonas, iavr}@image.ntua.gr

ABSTRACT

This paper describes a comprehensive framework giving support to a wide range of personalization facilities in a multimedia content management environment. The framework builds upon a rich, ontology-based representation of the domain of discourse, whereby content semantics are linked to a rich representation of user preferences. The expressive power of ontologies is used to develop automatic learning capabilities, in order to update user profiles as users interact with the system. The resulting descriptions of user interests in terms of ontologies are exploited, along with available content metadata, to provide users with personalized content search, browsing, ranking, and retrieval. On a wider perspective, the framework is built as an open platform that provides for further user and device adaptive capability extensions.

KEYWORDS

Personalization, user modeling, ontology, semantic web, information retrieval, multimedia

1. INTRODUCTION

The emergent semantic-based knowledge technologies [Berners-Lee 2001], also labelled "the semantic web" – a term which is falling short today for the breadth and reach of this area – aim at endowing software systems with a deeper insight on the meaning of the data they manipulate, create, store, and exchange. Functional areas where the capabilities enabled by such advanced semantic representations are being envisioned include, among others, sharing (exchange and integration), retrieval (search, filtering, browsing, recommendation), and presentation (visualization, navigation, composition) of application-domain knowledge. More recently, a growing interest is being raised for the potential of semantically rich descriptions to achieve improvements in the area of personalization technologies [Gauch 2004].

In the context of software systems and applications, personalization refers, in general terms, to the development of models and systems that represent and capture user preferences, goals, needs, knowledge, demographic information, environment, device, mood, capabilities, disabilities, etc., and use this information in the system to better meet user needs and expectations, and help users achieve their tasks and goals more efficiently [Kobsa 2001]. Personalization technologies gained significance in the 90's, with the boost of large-scale computing networks which enabled the deployment of services to massive, heterogeneous, and less predictable end-consumer audiences. As the number of services and the volume of content (text and multimedia; public, commercial and personal) in these networks keeps growing, personalization is more than ever a critical enabler in helping consumers manage capacity and complexity, and help vendors (content providers, managers, brokers, distributors, technology providers) reach their target audience and attain a competitive edge.

Semantic-based techniques enable to infuse software systems with a more precise understanding of application-domain knowledge, and henceforth, provide better means to define user needs, preferences, and activities within or with regards to the system. Moreover, they can be used for a richer representation of user-related information itself. In this paper we describe a personalization system that has been developed in this perspective. The system is part of a wider framework being developed in the aceMedia project [Kompatsiaris 2004]. aceMedia aims at integrating knowledge, semantics and multimedia technologies to solve user problems via intelligent content and applications. Using automatic content analysis tools, the aceMedia system augments multimedia items with self-descriptive metadata, thereby building up an understanding of the meaning of contents along different dimensions: application-domain concepts, visual semantics, media properties, formats, etc. aceMedia exploits this knowledge to provide new or better content services, such as intelligent search, or advanced browsing and navigation facilities [Grira 2004].

The goal of the aceMedia personalization system, presented here, is to exploit the semantic insights of the aceMedia framework to provide a comprehensive support for the development of a layer of semantic-aware user and device-adaptive capabilities. The current achievements in this undertaking include:

- The definition of a user profile representation that takes advantage of the expressive power of the ontologies and metadata that are used to describe contents. The model is prepared to comprehend a wide range of user-related information, in particular, user preferences for domain concepts, which are used to predict user preferences for contents annotated (automatically, by aceMedia content analysis tools) with such concepts.
- The development of automatic learning techniques to update user preferences by observing and analyzing user behavior, as recorded in histories of user actions within the system. These techniques can be complemented or supervised by the user herself, who can input new profile data manually, or correct the data inferred by the automatic update procedures.
- The development of a user profile management module to store, manage and provide access to the information captured in user profiles. This module makes information available, in the appropriate format, to any external module that may exploit it to improve system functionalities, such as intelligent retrieval, browsing, recommendation, or content adaptation, that are being developed in the aceMedia project.
- The definition of basic algorithms and measures to support preference-based content filtering and retrieval, making the most of the semantic representation of user preferences and the available ontology-based content metadata. This includes functions providing personalized content rating values, and personalized ranking methods to sort content sets while browsing, searching, or managing content collections.

The next section provides some preliminary context for our work, concerning when, how, and to what extent it should be appropriate to personalize in a system like aceMedia. An overview of the personalization system within the overall aceMedia framework is also given at the end of this section. Section 3 introduces our approach to the representation of user profiles. Section 4 explains the automatic update strategies for semantic user preferences, and Section 5 describes the exploitation of semantic preferences for personalized content retrieval. The paper ends with some conclusions and perspectives for future work.

2. PERSONALIZATION REQUIREMENTS

The automatic adaptation to expected user needs is not an easy task. It is sometimes appropriate to personalize system behavior, but at times personalization becomes obtrusive. Most of the time, personalization has to co-exist with other driving factors and priorities, so that the intensity and scope of automatic system adaptation has to be selected and balanced with great care.

Besides the purely innovative motivation, related to open problems and objectives of interest in the area of personalization, our work has been driven by a formal analysis of user requirements undertaken in the aceMedia project. This analysis was based on the elaboration of a set of context scenarios that, among many other aspects of the overall system, illustrate the envisioned personalization facilities. Using these scenarios, interviews with prospective system users gave rise to a large set of formal user requirements. In this section we summarize the main considerations issued from the analysis of these requirements.

2.1 When to personalize

The possibilities for the personalization of software systems are manifold. We have selected the ones that have been regarded as more valuable in user requirements and scenarios. The personalization features and functionalities that have been elicited as user-perceived qualities in user requirements and scenarios include learning preferences, preference-based content filtering, ranking, sharing, and recommendation. The identified situations where personalization is expected to show its benefit can be characterized as the following:

- *Imprecise user needs*. When the user makes a vague request, such as "select the best k items out of the N available", objective quality-based criteria may not suffice to provide a small enough selection. In situations like this, personalization can supply automatic criteria for filtering, sorting, and presenting contents based on an approximation to user preferences, for lack of explicit ones. This situation can be extended to ones where the user does input a query, but a very broad one. For example, a wide query concept such as "Sports" leaves more room for personalization than a narrower one like "Björn Borg".
- *Large amount of available contents*. Often more content items are available than the user can ever go through, even after the user has cut down the number with queries and conditions. In many cases, the filtering conditions, if any, do not provide a sufficient discrimination between the relevant contents, either because filters are boolean (e.g. the user is browsing by category and finds hundreds of items under "Formula 1", or is browsing by date and finds hundreds of items in the time range of interest), or the search rank values are very close to each other. At this point, a personalized preference-based rank can be used to scatter and sort results.
- *Short available user time, effort / quality tradeoffs*. The user could achieve the best results manually or by doing successive incremental searches, but due to time constraints or other priorities, he is willing to relax his goals in favor of quicker and easier to achieve results.
- *Initializing a search session*. Many intelligent iterative retrieval techniques start from a "page zero". Personalization can provide a focused initial content set, or even be intertwined at the request of the user at different stages of an incremental search (the last step being a particular case).
- *New content available*. The user is not aware of new incoming contents of potential interest in some repository. Not knowing what is available, she would need to repeat all her former queries to keep up to date. In general (take for example daily TV broadcasting), it is very hard to keep up to date with the pace and volume of content production this way, and personalization provides a sensible alternative. The notification updates can be fired on the initiative of the user herself ("Show me what's new"), of any other player in the production chain (an acquaintance wishing to share personal contents with friends a provider willing to advertise new items to potential customers, etc.), or automatically by the system.

The situations described above often overlap (e.g. the use has short time, imprecise needs, and too many contents are available), making personalization an all the more convenient feature.

Regarding non-functional requirements, privacy and reliability were two major and well-known concerns in the area of personalization technologies, also elicited in our user requirement analysis. Besides protecting (in terms of access and use) any sensitive information about the user that is stored in the system, and granting that usage data shall be recorded only under user consent, our system grants the right to view and edit any recorded or inferred information about the user. Since user preferences are quite complex in many respects, personalization technologies provide in general "good guesses", but do not (cannot) aim to provide perfect answers. Therefore, personalization effects are carefully tuned in our system, circumscribed to side options whenever appropriate, and providing the user with easy means to activate or deactivate personalization at any time. The user may amend wrong system guesses by inspecting and editing the automatically inferred preferences.

2.2 Hypothesis and principles

The initial hypothesis from which our research and development work on personalization starts off can be summarized as follows:

- Specific functionalities of the system where personalization effects would be perceived as a useful feature include (but are not limited to): intelligent search and retrieval, browsing and navigation, creation of content collections, albums, shows, and content sharing.

- Personalization can improve the precision (relevant items returned by search / total returned items) and recall (relevant items returned by search / total relevant items in the search space) of content retrieval (browsing and intelligent search), as subjectively perceived by users. The degree of improvement should:
  - Degrade monotonically with the specificity of user requests.
  - Increase monotonically with the number of results returned by non-personalized retrieval.
  - Decrease monotonically with the size of the rank values computed by non-personalized retrieval.
  - Increase monotonically with the concentration (i.e. low scattering) of non-personalized rank values (e.g. personalization would achieve the best improvements when retrieval is boolean, i.e. all positive results are equally ranked 1, and negative ones 0).
- Personalized content retrieval can achieve better search performance levels (e.g. precision and recall) than non-personalized retrieval at low (fixed) user-cost levels in terms of required effort (e.g. number of actions) and time from the user.
- As a complementary feature to content browsing and relevance feedback, personalization can speed up the retrieval sessions by providing focused content sets at different points in the session, at the user request.
- The less aware is the user of available contents, and the time when they are made available, the better enhancements can be attained by personalization.

Figure 1 illustrates some of the hypothesis above. Note that the double-line curve in the graphics represents the performance of personalized retrieval as perceived by the user (i.e. with respect to her subjective assessment of what contents are relevant). Whenever appropriate, this performance should be understood as that of personalization measures combined with retrieval measures (e.g. relevance for a query) computed by intelligent retrieval, not personalization alone (otherwise, for instance, in the top-right graphic, personalized retrieval performance should be constant, i.e. a straight horizontal line).

Figure 1. Hypothesis for personalized content retrieval
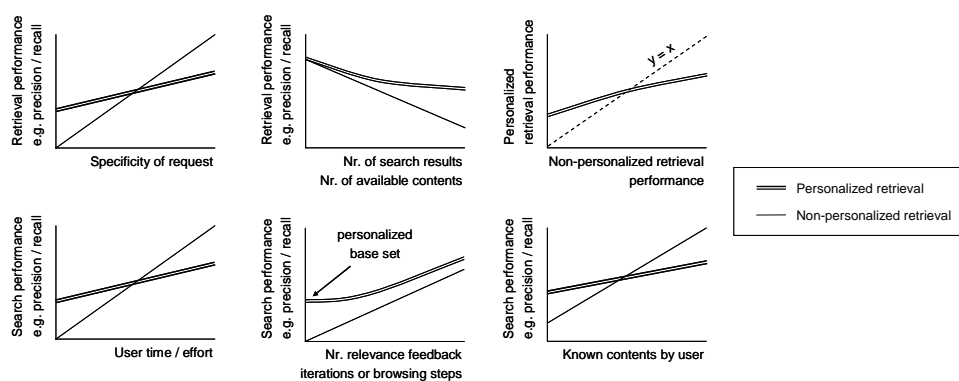
Our research with regards to the above hypothesis is twofold: provide the means to measure, validate, reject, or restate these hypothesis; and, to the extent that the hypothesis prove valid, identify and anticipate personalization failures and worst-case scenarios, in order to avoid or minimize the effects of expected failures.

2.3 Overview of the personalization framework

The global goal of our personalization framework is to provide comprehensive support for user awareness in a knowledge-driven multimedia management system. This involves representing and capturing all the relevant information about the user and her computing environment. The framework also facilitates the exploitation of this information for content retrieval, by providing ready-to-use algorithms to rate and rank contents according to user preferences.

In specific terms, the work achieved so far includes the definition of the knowledge structures for the representation of user profiles, including semantic preferences, media preferences, personal data, and device profile; the definition of algorithms for the automatic update of semantic user preferences by analyzing usage history; the definition of algorithms that provide personalized relevance measures and personalized ranking functions, for content retrieval (searching and browsing contents).

Figure 2. High-level view of the personalization system.
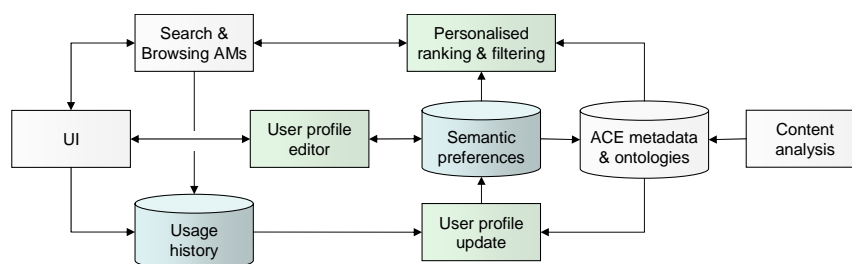
The personalization work is linked to several other components of the overall system (see Figure 2). It provides the search and browsing modules with algorithms for the personalization of content retrieval. These algorithms exploit content metadata automatically created by content analysis tools (or manually defined by users), and uses the semantic concept definitions provided by domain ontologies. So do the automatic user preference learning methods, which in addition, take as input the usage history collected from other system components, such as the browsing and searching modules, and the user interfaces.

3. USER PROFILE REPRESENTATION

A wide variety of user information has been modeled in different state-of-the-art personalization techniques. The selected set in aceMedia is driven by the user requirements to be met, the approach to meet them, the research goals to be achieved, and system requirements from other components external to the person-alization framework, that may need to store and access data or information in any way related to the user.

3.1 User profile contents

In our framework, user profiles include both static and dynamic information, information that needs to be entered by the user (e.g. personal data, media preferences), and information that is automatically acquired by the system (e.g. semantic user preferences, device capabilities). All this information is accessible to the user for inspection and editing at any time. User profiles are mainly explicit, although the usage history that is kept within the system can be seen as a form of implicit user model. In general, the personalization facilities do not directly access the usage history. Only the user profile update module does, in order to extract a high-level synthesis, more suitable for exploitation, in the form of semantic user preferences.

Semantic user preferences are the most complex part of user profiles. They exploit knowledge of domain concept semantics, and provide deep insights about user interests for these concepts and their properties. This is a rich part of user profiles in the sense that it enables powerful content personalization features. Besides the representation of semantic preferences, user profiles are open to extension to represent further user-related information. Currently, the definition of user profiles includes: device profile, media preferences, and personal data (name, birth date, nationality, residence, language, education, job). The current version of the personalization system only supports profiles for individual users. Group profiles could be addressed in the future. Of course, nothing prevents several users from sharing a common individual profile (e.g. by sharing a common login), but this would not provide any specifically group-aware personalization feature.

The formats for user profiles and usage history are based on XML. Device profiles and media user preferences are encoded in MPEG-21, to facilitate the communication with a cross-media engine able to automatically adapt media properties to different device capabilities and user preferences [Martínez 2004]. In future versions, the standards for ontology definition such as RDF or OWL might be considered for some parts of user profiles, with a view to facilitate the development of inference mechanisms on top of user profiles, using existing ontology-based reasoners.

3.2 Semantic preferences

In order to offer automatic extraction of user's semantic preferences and personalized enhanced search and retrieval capabilities to users of multimedia content, one technological target of is to offer semantic views to

existing a/v content, through the use of core and domain ontologies and to personalize those views according to the context defined by the profile of individual users or specific user groups; the latter clearly appreciating that semantic interpretation heavily relies on the context which in turn depends on the specific profile.

Semantic user preferences are represented by pairs concept-weight, where a concept is an instance or a class of the aceMedia domain ontology and the weight between -1 to 1 where a negative value indicates a negative preference of the user for the concept, and a positive value indicates a positive preference of the user for the concept. The whole set of preferences are clustered in a set of preferences scopes, either created by the user or inferred automatically by the user profile update module. These scopes could be for example a leisure scope, a work scope, travel scope etc.

The target of the personalization subtask can be identified as the implementation of the mechanism that creates, updates and uses user profiles, so that it acts as a personalizing agent residing between the end user and the content. The main principles followed for this development may be summarized in the following: The models are designed in a way that facilitates both the process of using user preferences in information retrieval as well as the process of tracking of usage history. Furthermore, they are designed in a way that allows for the automated extraction of user preferences from the usage history. Since several system modules (user profile storage, update, exploitation) depend on a common way of representing / exchanging data related to user profiles, an XML-based representation has been adopted to satisfy the requirements of the representation language for the semantic personalization structures.

4.   AUTOMATIC USER PROFILE UPDATE

The process that is responsible for the update of user preferences is executed in the offline part of the application. This means that its algorithms are not optimized to run in real-time. Since the process is not required to terminate in real-time, it is designed in a way to support the update of user preferences not based solely on the changes of usage history (incremental update).

More specifically, the profile update module receives as input usage history records of the system's end users and produces as output the corresponding weighted user preferences. In order to achieve this, the process also accesses information provided by the core and domain ontologies, as well as a custom-built semantic index, containing the semantic annotation of documents in the usage history and a custom knowledge model, based on the notion of a fuzzy relation. All information contained in usage history, the index and the knowledge model is used for the determination of weighted user preferences for each one of the system's domain concepts. In our current model, user queries are the part of usage history that dominates this analysis.

The main capability of the implemented approach is clustering of documents based on the semantic index, i.e., detecting which concepts are related to which document and then applying clustering algorithms to identify the concepts in the selected documents. Information is gathered from available users' semantic search actions and according to its kind necessary adaptations are performed, e.g. in the case of terms, mapping of terms and concepts is required as an intermediate step. Based on the notions of document clustering, the adopted methodology prepares the system's input for the profile update algorithms. Thus, the same approach is followed for personalization purposes and can be used for content categorization and manipulation.

The design of the user profile update module is based on previous expertise acquired in the field. The adopted knowledge model, together with the mathematical approach in the form of the semantic index, relies on the theory of fuzzy sets and relations [Klir 1995], implementing additional capabilities and providing added value to the corresponding personalization tasks, such as fuzzy in comparison to crisp interpretations and relationships between concepts at the semantic level. Preferences for semantic concepts are extracted by applying clustering algorithms [Miyamoto 1990] on the usage information data. While positive preferences extracted this way may be automatically recorded in the user profile, negative preferences first need to be validated in order to make sure that no information that the user might be interested in will be kept from him in the future. Therefore, for each one of the mined negative preferences a check is performed, ensuring that it does not contradict with any of the positive preferences that the profile contains. Since the mined sets have a great possibility of baring only a statistical meaning (rather than a semantic one), semantic cluster validation algorithms are also developed. In all cases of extraction of preferences, query keywords that were not preferred by the user are validated against positive preferences.

5. USER PROFILE EXPLOITATION

Exploiting user profiles involves using the information contained in profiles to adapt different features of the aceMedia system to particular user and device characteristics. The goals addressed so far in this part of the personalization system has been focused on delivering preference-based improvements for content filtering and retrieval, in a way that can be very easily introduced as an additional layer to system modules that support the retrieval functionalities, such as searching, browsing, and recommendation.

A fundamental notion for this purpose is the definition of a measure of content relevance for the interests of a particular user. Building on this measure, specific algorithms have been developed for filtering and sorting sets of content items (corresponding to entire content collections, categories being browsed, or base content sets to initialize relevance feedback), and re-ranking search results, according to user preferences.

In contrast with other previous approaches in personalized content retrieval, in our approach user profiles are explicit (as opposed to e.g. sets of preferred documents), consisting of weighted preferences for ontology concepts, rather that plain keywords or categories. This brings all the advantages of ontology-based representations, such as reducing ambiguities in the representation of user preferences, exploiting structural information of class hierarchies, or analyzing explicit relations of preferred concepts with other concepts in order to provide enhanced, semantic-aware personalization capabilities. Even more expressive user preferences could be considered as consisting not only of weighted ontology concepts, but ontology-based statements, restrictions, or logical expressions of arbitrary complexity. These preferences could be applied with the aid of ontology query engines and reasoners, but this is currently future work.

Figure 3. First system prototype
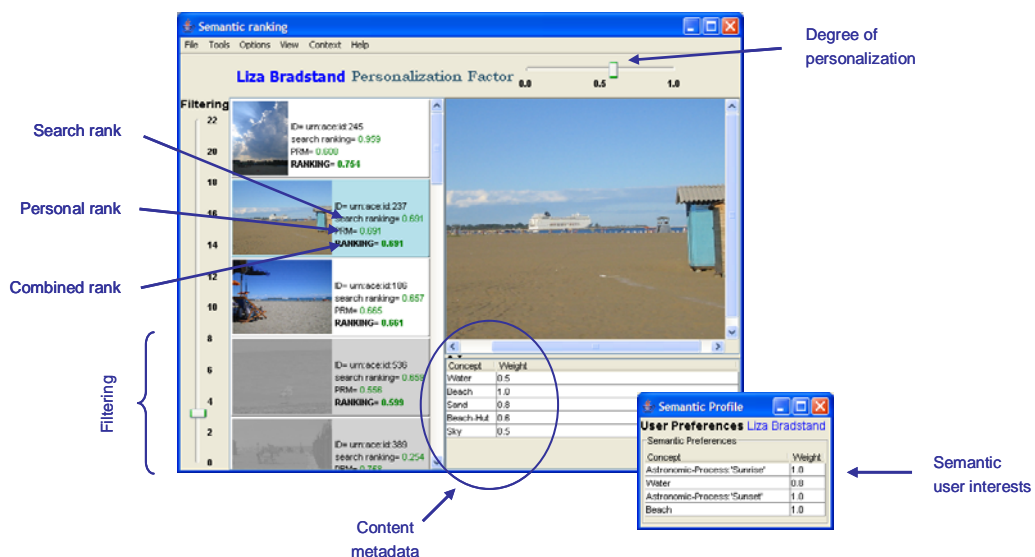
The interaction between the personalization module and the modules that will ultimately show the personalization effects (or provide the information to be personalized) is quite clean, since the information being exchanged consists of user IDs, lists of content item IDs, lists of content rank values, and optionally a parameter to gauge the degree of personalization. As a matter of fact, the personalization framework provides the core components and functionalities to support additional application features to those being currently tested, such as profile-based content sharing and recommendation, or automatic video summarization.

At the time of this writing, our techniques are being tested and evaluated on a first prototype implementation. Figure 3 shows a snapshot of a part of this prototype. The main window shows a list of ranked contents (in this example, still images) from a search result. When the personalized search option is turned on, the ranking changes to match the selected user preferences. Different user profiles can be selected, and the ordering changes when a different profile is selected.

As explained in the introductory sections, automatic personalization is not appropriate in all situations. Therefore it is considered an optional feature that users can turn on and off at any time. As a matter of fact, in the current prototype, a slider in the main window allows the user to set the degree of personalization as a

free parameter. As this parameter is moved between zero and total personalization, the list ordering moves dynamically, from the original non-personalized ranking, to maximum personalization. This is done according to a combination of the personal relevance measure and the search ranking. The UI also lets the user view the user's preferences and the metadata of content items.

6. CONCLUSION

The overall goal of the personalization framework described in this paper is to represent, capture, and manage any information about a multimedia management system users that can be exploited to improve the overall system services in different ways. The identification of the opportunities for improvement by personalization, and their achievement or facilitation, are part of the goals of this research.

The achievements reached so far include: an approach to the representation of semantic user preferences with respect to ontology-based content semantics; methods and algorithms for automatically learning user preferences; algorithms to measure personal content relevance based on ontology-based representations of content semantics and user preferences; algorithms for the personalized ranking of the results of intelligent search and retrieval; a first implementation of such methods, algorithms, and models, plus the components required for storage and management of the models, and the additional capabilities needed for the communication with the aceMedia framework and external application modules.

The work has been driven by the user requirement analysis conducted in the aceMedia project. The implementation of the personalization framework is currently being tested, extended and improved. The system has been designed as an open, modular framework, easy to extend or reconfigure, which, in particular, shall facilitate such progress and experimentation. The continuation of the work includes, among other tasks, refining and tuning the models, methods and algorithms developed so far, by testing and evaluating the framework; research and development of techniques for the automatic adjustment of personalized ranking algorithms; the adoption of a more powerful semantic, ontology-based, retrieval model [Vallet 2005]; the development of a user profile editor; the adoption of a privacy model and policies; and the extension of the framework to a distributed architecture, with mobile clients.

ACKNOWLEDGEMENT

REFERENCES

Berners-Lee, T. et al, 2001. The Semantic Web. *Scientific American*, Vol. 284, No. 5, pp. 34-43.

Gauch, S. et al, 2004. Ontology-Based Personalized Search and Browsing. *Web Intelligence and Agent Systems*, Vol. 1, No. 3-4, pp. 219-234.

Grira, N., Crucianu, M., Boujemaa, N., 2004. Fuzzy Clustering with Pairwise Constraints for Knowledge-Driven Image Categorization. *European Workshop on the Integration of Knowledge, Semantics and Digital Media Technology (EWIMT)*. London, U.K.

Klir, G. and Bo Yuan, 1995. *Fuzzy Sets and Fuzzy Logic: Theory and Applications*. Prentice Hall, New Jersey.

Kobsa, A., 2001. Generic User Modelling Systems. *User Modelling and User-Adapted Interaction*, Vol. 11, pp. 49-63.

Kompatsiaris, I., Avrithis, Y., Hobson, P., Strintzis, M. G., 2004. Integrating Knowledge, Semantics and Content for User-Centred Intelligent Media Services: the aceMedia Project. *Proceedings of 5th International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS)*. Lisbon, Portugal.

Martínez, J. M., Bescós, J. et al, 2004. Integrating Metadata-Driven Content Adaptation Approaches. *European Workshop on the Integration of Knowledge, Semantics and Digital Media Technology (EWIMT)*. London, U.K.

Miyamoto, S., 1990. *Fuzzy Sets in Information Retrieval and Cluster Analysis*. Kluwer Academic Publishers, Dordrecht.

Vallet, D., Fernández, M., Castells, P., 2005. An Ontology-Based Information Retrieval Model. *2nd European Semantic Web Conference (ESWC)*. Heraklion, Greece. Springer Verlag Lecture Notes in Computer Science.