A contextual personalization approach based on ontological knowledge

David Vallet[1] and **Miriam Fernandez**[2] and **Pablo Castells**[3] and **Phivos Mylonas**[4] and **Yannis Avrithis**[5]

Abstract. Combining traditional personalization techniques with novel knowledge representation paradigms, such as the ontology-based approach proposed in the Semantic Web field, is a challenging task. Personalization is a difficult problem when dealing with multimedia content and information retrieval, where context is increasingly acknowledged to be a key notion in order to make proper sense of user needs. This work focuses on contextualization within personalization in a multimedia environment. Towards that scope, we propose a novel contextual knowledge modeling scheme, and an approach for the dynamic, contextual activation of semantic user preferences to better represent user interests in coherence with ongoing user activities, e.g. in an interactive retrieval process. The application of this methodology is demonstrated using two user scenarios, and the performance results of a preliminary experiment are shown.

1 INTRODUCTION

Over the last decades, the task of personalization is related to various scientific and applied fields, with applications of techniques ranging from artificial intelligence and pattern recognition to traditional or multimedia databases and information retrieval applications [2]. One of the main issues arising is the problem of information overload, especially in the case of information retrieval that tends to select numerous multimedia documents, many of which are barely related to the user's wish [3]. This leads to other sources of information about user wishes and personalization is an approach that uses information stored in user preferences, additionally to the queries, to estimate the users' wishes and select the set of relevant documents.

In order to provide effective personalization techniques and develop intelligent personalization algorithms, it is appropriate not only to consider each user's queries/searches in an isolated manner, but also to take into account the surrounding contextual information available from prior sets of user actions. As an example, consider having some irregularities occurring in random places within a user's preferences, due to spontaneous changes of user's attention and focus. Taking into account further contextual information, the system can provide an undisturbed, clear view of the actual user's preferences, cleaned from extraordinary - according to each user's profiling information - anomalies, distractions or "noise" preferences. We refer to this surrounding information as contextual knowledge or just context.

Since several forms of context exist in the area [7], the problems to be addressed include how to represent context, how to determine it, and how to use it to influence the results of personalization. The idea behind the use of contextual information responds to the fact that not all human acts are relevant in all situations and since context is a difficult notion to grasp and capture, we restrict it herein to the notion of *ontological context*. The latter is defined as a "fuzzified" version of traditional ontologies [5]. This work is concerned with exploiting semantic, ontology-based contextual information aimed towards its use in personalization tasks. The effect and utility of the proposed invention consists of endowing a personalized retrieval system with the capability to filter and focus its knowledge about user preferences on the semantic context of ongoing user activities, so as to achieve a coherence with the thematic scope of user actions at runtime. The difficulty of successfully applying extraction of user preferences in multimedia environments, using an ontological knowledge representation constitutes this task an open and challenging issue. Finally, in the context of the Semantic Web, research efforts have resulted in the development of new knowledge representation languages, such as RDF, utilized throughout the current approach.

The rest of the paper is organized as follows: in section 2, we present the main components of the underlying knowledge infrastructure, introducing the notion of fuzzified context, as well as the use of fuzzy relations within ontologies. Section 3 deals with the problems of runtime context determination and context usage in order to influence activation of user preferences, "contextualize" them and predict or take into account the drift of preferences over time. As will be described a runtime context is represented as a set of weighted concepts from the domain ontology. How this set is determined, updated, and interpreted, will also be explained. In section 4 we provide early experimental results in the form of two user case-study examples and some conclusions are drawn in section 5.

2 ONTOLOGY-BASED KNOWLEDGE REPRESENTATION

Knowledge representation is one of the central and in some ways most fundamental notions in fields like information retrieval. Different views have been proposed and studied, and attempts have been made at determining what representation properties are important for knowledge representation in multimedia applications. However, most proposed solutions are not sufficient due to performance reasons, as well as due to the lack of accompanying contextual information. The latter forms a major limitation and it lies within the intensions of current work to manipulate and improve this kind of infor-

[1] Universidad Autonoma de Madrid, Escuela Politecnica Superior, Spain, email: david.vallet@uam.es
[2] Universidad Autonoma de Madrid, Escuela Politecnica Superior, Spain, email: miriam.fernandez@uam.es
[3] Universidad Autonoma de Madrid, Escuela Politecnica Superior, Spain, email: pablo.castells@uam.es
[4] National Technical University of Athens, Image, Video & Multimedia Laboratory, Athens, Greece, email: fmylonas@image.ntua.gr
[5] National Technical University of Athens, Image, Video & Multimedia Laboratory, Athens, Greece, email: iavr@image.ntua.gr

mation in an efficient manner. Unquestionably, design and analysis of such a task is not straight-forward and many approaches are acceptable. The term context can take many interpretations and definitions when dealing with specific application-domains [7]. This statement denotes the need for a working context interpretation applicable in personalization, since both domains will benefit from and contribute to each other. A restriction of the general notion of context is necessary, identifying the type of context suitable for user profiling and extraction of user preferences. This kind of context is defined with the aid of fuzzy algebra and ontologies, as a "fuzzified" version of traditional ontologies. We shall use the term *ontological context* from now on.

An ontology is a formal specification of a shared understanding of a domain [5]. This formal specification is usually carried out using a subclass hierarchy with relationships among the classes, where one can define complex class descriptions (e.g. in in DL [1] or OWL [8]), and use a reasoner to infer new relations among ontology elements. Given a specific domain $\mathcal{O}$ and using relations $\mathcal{R}$ and appropriate semantics, an ontology can be modeled as a set of concepts $\mathcal{C}$ together with the corresponding relations $\mathcal{R}$ between the concepts of the domain: $\mathcal{O} = \{\mathcal{C}, \{\mathcal{R}_i\}\},\ \mathcal{R}_i : \mathcal{C} \times \mathcal{C} \to \{0,1\},\ i = 1 \ldots n,\ n \in \mathbb{N}$. In this formula, $\mathcal{R}_i$ denotes the i-th relation between the concepts in the ontology.

Although in general any type of relations may be taken into consideration, in order to extract and use the desired ontological context, we define it in the means of *fuzzy ontological relations*. Fuzziness is an intrinsic property of knowledge representation, since accurate representation of real-life information is only achieved through the use of fuzzy relations. In [13] a set of basic relations is proposed that can be used to model taxonomic context hierarchies, while the relations themselves represent deeper semantics than just a taxonomic relation. Without claiming that the proposed relations are sufficient to model every type of context, we think that the relations presented in Table 1 are generic enough to form a useful basis for our personalized context model.

Table 1. Ontological relations suitable for personalization

Abbreviation	Name	Description
$Pr(x,y)$	PropertyOf	x is the property of y
$P(x,y)$	PartOf	x is part of y
$Sp(x,y)$	SpecializationOf	x is specialization of y, i.e. this corresponds to the well-know subclass relation
$Ct(x,y)$	ContextOF	x provides the context for y
$Loc(x,y)$	LocationOf	x is the location of y
$Pr(x,y)$	PropertyOf	x is the property of y

The presented relations are based on the set of semantic relations defined by the MPEG-7 standard [12]. Consequently, we may fuzzify the previous formula and describe an ontology suitable for personalization by using the following notation: $\mathcal{O}_{\mathcal{F}} = \{\mathcal{C}, \{\mathcal{R}^{\mathcal{F}}_{c_i,c_j}\}\},\ \mathcal{R}^{\mathcal{F}}_{c_i,c_j} : \mathcal{C} \times \mathcal{C} \to [0,1],\ i,j = 1 \ldots n,\ n \in \mathbb{N},\ i \neq j$. This context model forms an ontology itself, as it is compatible with the above definition. We use this "fuzzified" definition of the knowledge model in the following sections of this paper, since it is considered to be the most suitable for the modeling of information governed by uncertainty and fuzzified relations, like in the real world.

Finally, when dealing with implementation issues of the proposed context knowledge representation, we propose a specific way of representing context, following a standardized language like OWL or RDF. We smoothly integrate context's functionalities in the ontology infrastructure, i.e. we adopt enhanced characteristics available in the area of the Semantic Web, like the reification technique [10]. The proposed context model is described by pairs of concepts, represented as ontology *classes*, and relationships between the pair members, represented by *properties*. To introduce fuzziness in the approach, a degree of confidence is attached to each property. Non-existing relationships between concepts imply non-existing fuzzy relations, i.e. relations with zero confidence values are omitted. Additionally, every concept participating in the contextualized ontology has a unary degree of confidence to itself, apart from the degrees of confidence that exist between any possible class interconnections.

3 CONTEXTUAL PERSONALIZATION

Having fulfilled the first step towards contextual personalization in the form of contextual knowledge representation, the next basic step to consider is the definition of a strategy on dynamic contextualization of user preferences. Three basic principles dominate the latter:

1. representation of context as a set of domain ontology concepts that a user has "touched" or followed in some manner,
2. extension of this representation of context by using explicit semantic relations among concepts represented in the ontology
3. extension of user preferences by a similar principle

Roughly speaking, the "intersection" of the above two sets of concepts, with combined weights, are taken as the user preferences. In the following, an approximation to conditional probabilities will be utilized as an ontology-based extension mechanism. The latter is based on the existence of relations between concepts. More formally, given a finite set Ω, and $\alpha \in \Omega$, let $P(\alpha)$ be the probability that α holds some condition. We shall use this form of estimating "the probability that α holds some condition" with the purpose of extending user preferences for ontology concepts. The condition will be "the user is interested in concept α", that is, $P(\alpha)$ will be interpreted as the probability that the user is interested in concept α of the ontology. Universe Ω will correspond to a domain ontology $\mathcal{O}$ (the universe of all concepts). In the process of preferences and context expansion, a variation of constrained spreading activation (CSA) strategy is utilized [4], [11].

3.1 Semantic context for personalized content retrieval

Our model for context-based personalization can be formalized as follows: let $\mathcal{U}$ be the set of all users, let $\mathcal{C}$ be the set of all contexts, and $\mathcal{P}$ the set of all possible user preferences. Since each user will have different preferences, let $P : \mathcal{U} \to \mathcal{P}$ map each user to his/her preference. Similarly, each user is related to a different context at any given time, which we represent by a mapping $C : \mathcal{U} \times \mathbb{N} \to \mathcal{C}$, since we assume that context evolves over time. Thus we shall often refer to the elements from $\mathcal{P}$ and $\mathcal{C}$ as in the form $P(u)$ and $C(u,t)$ respectively, where $u \in \mathcal{U}$ and $t \in \mathbb{N}$. We define the contextualization of preferences as a mapping $\Phi : \mathcal{P} \times \mathcal{C} \to \mathcal{P}$ so that for all $p \in \mathcal{P}$ and $c \in \mathcal{C}$, $p| = \Phi(p,c)$.

In this context the entailment $p| = q$ means that any consequence that could be inferred from q could also be inferred from p. For instance, given a user $u \in \mathcal{U}$, if $P(u) = q$ implies that u "likes x" (whatever this means), then u would also "like x" if her/his preference was p. Now we can particularize the above definition for a specific representation of preference and context. In our model, we consider user preferences as the weighted set of domain ontology concepts for which the user has an interest, where the intensity of interest can range from 0 to 1. Given the domain ontology $\mathcal{O}$, we define the set of all preferences over $\mathcal{O}$ as $P_{\mathcal{O}} = [0,1]^{|\mathcal{O}|}$, where given $p \in \mathcal{P}_{\mathcal{O}}$, the value p_x represents the preference intensity for a concept $x \in \mathcal{O}$ in the ontology. Under the above definitions, we particularize $| =_{\mathcal{O}}$ as follows: given $p, q \in \mathcal{P}_{\mathcal{O}}, p| =_{\mathcal{O}} q \Leftrightarrow \forall x \in \mathcal{O}$, either $q_x \leq p_x$, or q_x can be deduced from p using consistent preference extension rules over $\mathcal{O}$. Additionally, we define the set of all semantic runtime contexts as $\mathcal{C}_{\mathcal{O}} = [0,1]^{|\mathcal{O}|}$. In the next sections, we propose a method to build the values of $C(u,t)$ during a user session, a model to define Φ, and the techniques to compute it. Once we define this, the activated user preferences in a given context are given by $\Phi(P(u), C(u,t))$.

3.2 Semantic extension of context

As already mentioned, the selective activation of user preferences is based on an approximation to conditional probabilities: given $x \in \mathcal{O}$ with $P_x(u) > 0$ for some $u \in \mathcal{U}$, i.e. a concept on which a user u has some interest, the probability that x is relevant for the context can be expressed in terms of the probability that x and each concept y directly related to x in the ontology belong to the same topic, and the probability that y is relevant for the context. With this formulation, the relevance of x for the context can be computed by a constrained spreading activation algorithm, starting with the initial set of context concepts defined by C.

Our strategy is based on weighting each semantic relation r in the ontology with a measure $w(r)$ that represents the probability that given the fact that $r(x,y)$, x and y belong to the same topic. We will use this as a criteria for estimating the certainty that y is relevant for the context if x is relevant for the context, i.e. $w(r)$ will be interpreted as the probability that a concept y is relevant for the current context if we know that a concept x is in the context, and $r(x,y)$ holds. Based on this measure, we use a constrained spreading activation strategy over the semantic network defined by semantic relations in the ontology, to expand the set of context concepts. As a result of this strategy, the initial context $C(t)$ is expanded to a larger context vector $EC(t)$, where of course $EC_x(t) \geq C_x(t)$ for all $x \in O$. Since $\mathcal{R}$ is the set of all relations in $\mathcal{O}$, let $\widehat{\mathcal{R}} = \mathcal{R} \bigcup \{r^{-1} | r \in \mathcal{R}\}$, and $w : \widehat{\mathcal{R}} \to [0,1]$. The extended context vector $EC(t)$ is computed by:

$$EC_y(t) = \begin{cases} C_y(t) \quad if \quad C_y(t) > 0 \\ R\left(\{EC_x(t) \cdot w(r) \cdot power(x)\}_{x \in \mathcal{O}, r \in \widehat{\mathcal{R}}, r(x,y)}\right) \end{cases}$$

where R is defined as:

$$R(X) = \sum_{S \subset \mathbb{N}_n} (-1)^{|S|+1} \prod_{i \in S} x_i$$

and $X = \{x_i\}_{i=0}^n$, where $x_i \in [0,1]$ and $power(x) \in [0,1]$ is a propagation power assigned to each concept x (by default, $power(x) = 1$).

3.3 Semantic preference expansion

In information retrieval two major issues need to be considered towards the efficient manipulation and exploitation of user preferences. The first thing to consider is their ability to adapt to the contextual environment, i.e. their context adaptiveness, and the second thing is the special care that needs to be taken for an overall profile consistency after application of user preferences contextualizing methodologies. Under these circumstances, a novel approach is followed: extension of preferences through ontology relations, following in general the same approach, that is used to expand the runtime context.

The idea behind this methodology is to follow the principles used for the extension of the semantic context in the previous subsection 3.2. The main difference is that here relations are assigned different weights $w'(r)$ for propagation, since the inferences one can make on user preferences, based on the semantic relations between concepts, are not necessarily the same as one would make for the contextual relevance. In general, it is expected that $w'(r) \leq w(r)$, i.e. user preferences are expected to have a shorter expansion than context has. Given an initial user preference P, the extended preference vector EP is defined by:

$$EP_y = \begin{cases} P_y, \quad if \quad P_y > 0 \\ R\left(\{EP_x \cdot w'(r) \cdot power(x)\}_{x \in \mathcal{O}, r \in \widehat{\mathcal{R}}, r(x,y)}\right) \quad otherwise \end{cases}$$

which is equivalent to the previous formula for $EC_y(t)$, where EC, C and w have been replaced by EP, P and w', and variable t has been removed, since long-term preferences are taken to be stable along small amounts of time.

3.4 Contextual activation of preferences

After expanding of context, only preferred concepts with a context value different from zero will count for personalization. This is done by computing a contextual preference vector CP, as defined by $CP_x = EP_x \cdot C_x$ for each $x \in \mathcal{O}$, where EP is the vector of extended user preferences. Now CP_x can be interpreted as a combined measure of the likelihood that concept x is preferred and how relevant the concept is to the current context. Note that this vector is in fact dependent on user and time, i.e. $CP(u,t)$. At this point we have achieved a contextual preference mapping as defined in subsection 3.1, namely $\Phi(P(u), C(u,t)) = CP(u,t)$, where $P(u)| = \Phi(P(u), C(u,t))$, since $CP_x(u,t) > P_x(u,t)$ only when $EP_x(u)$ has been derived from $P(u)$ through the constrained spreading expansion mechanism, and $CP_x(u,t) < EP_x(u)$.

4 EXPERIMENTAL RESULTS

The contextualization techniques described in this work have been implemented within an experimental prototype. We have tested the prototype on a medium-scale corpus in order to put to trial the validity and soundness of the proposed model, tune parameters, observe the behavior of the contextualization system, and draw some empirical results. In subsection 4.1 we present an example scenario of detecting a user's preferences and in subsection 4.2 we provide some early evaluation results of our methodology.

4.1 An example scenario

An example scenario is provided in this section as an illustration of the application of the contextual personalization techniques. For the sake of clarity and space, full account of example details, such as the entire set of ontology concepts and relations involved, are omitted. Let us consider that Clio's family and friends have set up a common repository where they upload and share their pictures. Clio has not checked out the collection for quite a while, and she is willing to take a look of what images her relatives have brought from their last summer vacations. Let us also assume that the proposed framework has learned some of Clio's preferences over time, i.e. Clio's profile includes the weighted semantic preferences for domain concepts of the ontology, shown in Table 2, where *Tobby* is her brother's pet and an instance of *Dog*:

Table 2. Clio's initial preferences

$P(Clio)$	
Car	1.0
City	1.0
Sea	1.0
Tobby	1.0
Vegetation	1.0

This would define the P vector. Now suppose that Clio selects two images shown in Figure 1. As a consequence, a runtime context is built including the elements shown in Table 3, which corresponds to the C vector.

Figure 1. Clio's selection of pictures

Table 3. Runtime context built

$C(Clio, 1)$	
Construction	1.0
Flower	1.0

Now, Clio wants to see some picture of her family members, and issues the query "my family". The contextualization mechanism comes into place and the context set is expanded through semantic relations from the initial context, adding two more weighted concepts, shown in bold in Table 4.

This makes up the EC vector. Similarly, Clio's preferences are extended through semantic relations from her initial ones. The expanded preferences stored in the EP vector are the following, where we show the new concepts in bold (Table 5).

The contextual preferences are computed by multiplying the coordinates of the EC and the EP vectors one-on-one, yielding the CP vector depicted in Table 6 (concepts with weight 0 are omitted).

Comparing this to the initial preferences in Clio's profile, we can see that *Car*, *Sea* and *Tobby* are disregarded as out of context preferences, whereas *Construction* and *Flower* have been added because they are strongly semantically related both to the initial Clio's preferences and to the current context.

4.2 Evaluation of contextual personalization

In general, evaluating personalization tasks is known to be a difficult and expensive task [14], [9]. We have conducted a preliminary ex-

Table 4. Clio's expanded context

$EC(Clio, 1)$	
Construction	1.0
City	**0.6**
Flower	1.0
Vegetation	**0.5**

Table 5. Clio's expanded preferences

$EP(Clio)$			
Car	1.0	***Tree***	**1.0**
City	1.0	***Road***	**0.5**
Construction	0.7	*Sea*	1.0
Dog	0.3	*Tobby*	1.0
Lake	0.8	*Vegetation*	1.0
Flower	1.0	***Water***	0.7
Plant	1.0		

perimentation of the proposed contextualization techniques, in order to test the feasibility, soundness, and technical validity of the defined models. To this end, we have set up a corpus of significant size consisting of 145,316 text documents (445MB) from the CNN web site, plus the KIM publicly available domain ontology and KB [6]. The KB contains a total of 281 RDF classes, 138 properties, 35,689 instances, and 465,848 sentences. The CNN documents are annotated with KB concepts, amounting to over three million annotation links. The relation weights were first set manually on an intuitive basis, and tuned empirically afterwards by running a few trials. In order to extract precision and recall figures, we have rated the document/query/preference/context tuples manually. Needless to say, this is by no means a valid evaluation, but rather a first step to check the consistency of the models, to debug and tune the functions and parameters and to make some preliminary observations.

Since the contextualization techniques are applied in the course of a user session, a sequence of steps needs to be defined in order to put them to work. According to this, we use again a short scenario, as follows: Clio is fond of all kinds of luxurious and stylish articles. Her preferences include fancy brands such as Rolex, Maybach, Lexus, Hilton, Aston Martin, Bentley, Louis Vuitton, Sony, Apple, Rolls-Royce, Mercedes, Ferrari, Prada, and BMW, among others. Clio starts a search session with a query for news about Daimler-Chrysler and the different brands the company owns. Daimler-Chrysler owns both luxury brands as Mercedes or Maybach, and other more ordinary ones like Dodge or Setra that are not of interest to Clio. Personalization reorders the results according to Clio's preferences, showing first the documents related to Daimler-Chrysler and its brands Mercedes or Mayback, and pushing down other documents related to the lower-end brands of the company. In consequence, person-

Table 6. Clio's contextual preferences

$CP(Clio, 1)$	
Construction	0.7
City	0.6
Flower	1.0
Vegetation	0.5

alized search performs significantly better from the user's point of view. Since this is the first query of the session, no context exists yet and at this point there is no measure of the performance of the contextualization techniques.

In order to obtain this kind of information, Clio opens some documents in the search result, about the Mercedes brand and how Daimler-Chrysler is going to commercialize a new car model. She also opens a document about the new model Maybach 62. The context monitor extracts the concept of Mercedes from the documents opened by the user, along with the concept Maybach, since the selected documents were mainly about these two brands. Next, Clio makes a new query: "companies that trade on the New York Stock Exchange and have brands in the USA". The context is expanded to new concepts such as Daimler-Chrysler, owner of Mercedes and Maybach, along with all its brands. The query results are resorted according to the contextualized preferences of Clio. The documents that mention Daimler-Crhysler and Mercedes are pushed up in the result set. Clio still encounters other companies and brands that trade in the New York stock exchange and match her preferences, like the Sony Corporation, but these are not found semantically close to the brands in the context, and therefore get a lower sorting that other contents more in context with the previous user actions, which explains the improvement shown in Figure 2.

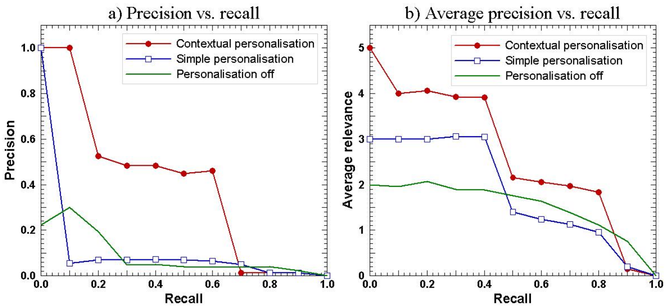

Figure 2. Comparative performance of personalized search with and without contextualization, for the query "Companies that trade on the New York Stock Exchange and commercialize a brand in the USA". The graphic a) shows the precision vs. recall curve, and b) shows the average relevance vs. recall.

5 CONCLUSIONS

Reliability is a well-known concern in the area of personalization, and one important source of inaccuracy of automatic personalization techniques is that they are typically applied out of context. I.e. although users may have stable and recurrent overall preferences, not all of their interests are relevant all the time. Instead, only a subset is usually active for them at a given time. What are the driving factors that determine this subset in a given situation is a hard question in general, all the more difficult to grasp and formalize in a computer system. Indeed, a wide range of procedural, cognitive, and environmental factors intervene in the dynamic orientation of user focus while s/he interacts with a system. The notion of context becomes elusive if one aims at a holistic approach. In this paper we propose an approximation to this problem on a specific perspective, namely based on a model of semantic runtime context of user actions, with the aim to achieve a perceivable improvement in the combination of personalization and content retrieval techniques.

As widely acknowledged, context is an increasingly common notion in information retrieval. In our approach, we combined traditional personalization techniques with novel knowledge representations, such as ontologies and reification. The use of semantic concepts, rather than plain terms, for the representation of contextual meanings, and exploitation of explicit ontology-based information attached to the concepts forms a significant novelty in the area. We also combined implicit context meanings collected at runtime, with a persistent and more general representation of user preferences. Benefit from the overall methodology is twofold: personalization techniques gain accuracy and reliability by avoiding the risk of having locally irrelevant user preferences getting in the way of a specific and focused user retrieval activity. Inversely, the pieces of meaning extracted from the context are filtered, directed, enriched, and made more coherent and meaningful by relating them to user preferences. Further insights to be drawn upon theoretic analysis, as well as the observations concerning performance and trade-offs from the experimental results of future implementation work and testing, shall provide further grounds for the analysis and evaluation of this approach.

REFERENCES

[1] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, *The Description Logic Hand-book: Theory, Implementation and Application*, Cambridge University Press, 2002.

[2] R.A. Baeza-Yates and B.A. Ribeiro-Neto, *Modern Information Retrieval*, ACM Press, Addison-Wesley, 1999.

[3] P.M. Chen and F.C. Kuo, *An information retrieval system based on a user profile*, Journal of Systems and Software 54, 3-8, 2000.

[4] F. Crestani, *Application of Spreading Activation Techniques in Information Retrieval*, Artificial Intelligence Review 11, 1997, pp. 453-482.

[5] T.R. Gruber, *A Translation Approach to Portable Ontology Specification*, Knowledge Acquisition 5: 199-220, 1993.

[6] A. Kiryakov, B. Popov, I. Terziev, D. Manov, and D. Ognyanoff, *Semantic Annotation, Indexing, and Retrieval*, Journal of Web Sematics 2(1), 2004, pp. 49-79.

[7] Ph. Mylonas and Y. Avrithis, *Context modeling for multimedia analysis and use*, Proc. of 5th International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT '05), 2005.

[8] W3C Recommendation, *OWL Web Ontology Language Reference*, http://www.w3.org/TR/owl-ref/, 10 February 2004.

[9] B. Rajagopalan and A. Deshmukh, *Evaluation of Online Personalization Systems: A Survey of Evaluation Schemes and A Knowledge-Based Approach*, Journal of Electronic Commerce Research 6(2), 2005, pp. 112-122.

[10] RDF Semantics, W3C Recommendation, *Reification, Containers, Collections and rdf:value*, http://www.w3.org/TR/rdf-mt/#ReifAndCont, 10 February 2004.

[11] G. Salton, C. Buckley, *On the Use of Spreading Activation Methods in Automatic Information Retrieval*, 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 98), Melbourne, Australia, 1988, pp. 147-160.

[12] T. Sikora, *The MPEG-7 Visual standard for content descritpion - an overview*, IEEE Trans. on Circuits and Systems for Video Technology, special issue on MPEG-7, 11(6):696-702, June 2001.

[13] M. Wallace, G. Akrivas, P. Mylonas, Y. Avrithis and S. Kollias, *Using Context and Fuzzy Relations to Interpret Multimedia Content*, Proc. of the 3rd International Workshop on Content-Based Multimedia Indexing (CBMI), IRISA, Rennes, France, September 2003.

[14] R. Wilkinson and M. Wu, *Evaluation Experiments and Experience from the Perspective of Interactive Information Retrieval*, 3rd Workshop on Empirical Evaluation of Adaptive Systems, in conjunction with the 2nd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH2004), Eindhoven, The Netherlands, 2004, pp. 221-230.