Personalized Content Retrieval in Context Using Ontological Knowledge

David Vallet, Pablo Castells, Miriam Fernández, Phivos Mylonas, and Yannis Avrithis

Abstract— **Personalized content retrieval aims at improving the retrieval process by taking into account the particular interests of individual users. However, not all user preferences are relevant in all situations. It is well known that human preferences are complex, multiple, heterogeneous, changing, even contradictory, and should be understood in context with the user goals and tasks at hand. In this paper we propose a method to build a dynamic representation of the semantic context of ongoing retrieval tasks, which is used to activate different subsets of user interests at runtime, in a way that out–of-context preferences are discarded. Our approach is based on an ontology-driven representation of the domain of discourse, providing enriched descriptions of the semantics involved in retrieval actions and preferences, and enabling the definition of effective means to relate preferences and context.**

Index Terms — **Content search and retrieval, Context modeling, Ontology, Personalization**

I. INTRODUCTION

THE size and the pace of growth of the world-wide body of available information in digital format (text and audiovisual) constitute a permanent challenge for content retrieval technologies. People have instant access to unprecedented inventories of multimedia content world-wide, readily available from their office, their living room, or the palm of their hand. In such environments, users would be helpless without the assistance of powerful searching and browsing tools to find their way through. In environments lacking a strong global organization (such as the open WWW), with decentralized content provision, dynamic networks, etc., query-based and browsing technologies often find their limits.

David Vallet, Pablo Castells, and Miriam Fernández are with the Universidad Autónoma de Madrid, Escuela Politécnica Superior, 28049 Madrid, Spain (phone: +34-914972284, fax: +34-914972235, e-mail: david.vallet@uam.es, miriam.fernandez@uam.es, pablo.castells@uam.es).

Phivos Mylonas and Yannis Avrithis are with the National Technical University of Athens, Image, Video & Multimedia Laboratory, 9, Iroon Polytechniou str., 15773 Zographou, Athens, Greece (email: fmylonas@image.ntua.gr, iavr@image.ntua.gr).

Personalized multimedia content access aims at enhancing the information retrieval (IR) process by complementing explicit user requests with implicit user preferences, to better meet individual user needs [15], [20]. Personalization is being currently envisioned as a major research trend to relieve information overload, since IR usually tends to select the same content for different users on the same query, many of which are barely related to the user's wish [9]. The combination of long-term and short-term user interests that takes place in a personalized interaction is delicate and must be handled with great care in order to preserve the effectiveness of the global retrieval support system, bringing to bear the differential aspects of individual users while avoiding distracting them away from their current specific goals.

Reliability is indeed a well-known concern in the areas of user modeling and personalization technologies. One important source of inaccuracy of automatic personalization techniques is that they are typically applied out of context. In other words, although users may have stable and recurrent overall preferences, not all of their interests are relevant all the time. Instead, usually only a subset is active at a given situation, and the rest can be considered as "noise" preferences. In order to provide effective personalization techniques and develop intelligent personalization algorithms, it is appropriate to not only 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.

It is common knowledge that several forms of context exist in the area [23]. This paper is concerned with exploiting semantic, ontology-based contextual information, specifically aimed towards its use in personalization for content access and retrieval. Among the possible knowledge representation formalisms, ontologies present a number of advantages [30], as they provide a formal framework for supporting explicit, machine-processable semantics definitions, and facilitate inference and derivation of new knowledge based on already existing knowledge.

The goal of the research presented herein is to endow personalized mutimedia management systems with the capability to filter and focus their 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. We propose a method to build a dynamic representation of the semantic context of ongoing retrieval tasks, which is used to activate different subsets of user interests at runtime, in such a way that out-of-context preferences are discarded.

Our approach is based on an ontology-driven representation of the domain of discourse, providing enriched descriptions of the semantics involved in retrieval actions and preferences, and enabling the definition of effective means to relate preferences and context.

The extraction and inclusion of real-time contextual information as a means to enhance the effectiveness and reliability of long-term personalization enables a more realistic approximation to the highly dynamic and contextual nature of user preferences, in a novel approach with respect to prior work. The gain in accuracy and expressiveness obtained from the ontology-based approach brings additional improvements in terms of retrieval performance. Furthermore, the semantic approach is key for the applicability of our symbolic methods to multimedia corpora, by relying on (manual or automatic [24]) semantic annotations of the content.

The rest of the paper is organized as follows. Section II introduces the notion of context and related work in this area. Our approach to contextual personalization is described in detail in Section III, including our underlying ontology-based personalization framework (Subsection III.A), the proposed context representation model (Subsection III.B), a mechanism to instantiate the model (Subsection III.C), a method to filter user preferences by context (Subsection III.D), and the final computation of a personalized retrieval function for preference-biased, context-sensitive result ranking (Subsection III.E). A detailed use case is provided in Section IV, and our initial experimental results are reported in Section V. Finally, some conclusions are given in Section VI.

II. THE NOTION OF CONTEXT

In order to address some of the limitations of classic personalization systems, researchers have looked to the new emerging area defined by the so-called context-aware systems [5]. In this scope, the term *context* can take on many meanings and there is not one definition that is felt to be globally satisfactory and that covers all the ways in which the term is used [12]. The term has a long history in diverse areas of computer science, namely in artificial intelligence, IR, image and video analysis, context-sensitive help, multitasking context switch, psychological contextual perception, and so on.

The effective use of context information in computing applications still remains an open and challenging problem. Several researchers have tried over the years to categorize context-aware applications and features, including contextual sensing, contextual adaptation, contextual resource discovery and contextual augmentation (the ability to associate digital data with a user's context) [25], [29]. These ideas can be combined and applied to the presentation of information and services to a user, the automatic execution of a service, or the tagging of context to information for later retrieval [1].

This paper is concerned with exploiting contextual information and smoothly integrating it into the personalization of content retrieval. In this field, contextual information can be proven to be very helpful when dealing with content retrieval queries and requests. Most existing IR systems base their retrieval decision solely on queries and document collections; information about actual users and search context is largely ignored, and as result a significant number of misclassifications occur.

Context-sensitive retrieval has been identified as a major challenge in IR research. Several context-sensitive retrieval algorithms exist in the literature, most of them based on statistical language models to combine the preceding queries and clicked document summaries with the current query, for better ranking of documents [3], [14], [16], [17], [21]. Towards the optimal retrieval system, the system should exploit as much additional contextual information as possible to improve the retrieval accuracy, whenever this is available [2]. One common solution is the use of relevance feedback [28]. However, the effectiveness of relevance feedback is considered to be limited in real systems, basically because users are often reluctant to provide the required information.

For this reason, implicit feedback has recently attracted greater attention [6], [18]. For a complex or difficult information request, the user may need to modify his/her query and view ranked documents in many iterations before the information need is satisfied. In such an interactive retrieval scenario, the information naturally available to the retrieval system is more than just the current user query and the document collection – in general, arbitrary interaction history can be made available to the retrieval system, including past queries, the documents that the user has chosen to view, and even how a user has accessed a document, e.g. via his/her Personal Digital Assistant (PDA) or Personal Computer (PC), in a read-only or read/write mode of usage, for how long, etc. Our research aims at enhancing the accuracy and effectiveness of prior approaches by a) using an enriched representation of the semantics of contents in the retrieval space, and b) combining information from the short-term retrieval context with a representation of longer-term user interests, to gain a subjective improvement for an individual searcher.

III. PERSONALIZATION IN CONTEXT: OUR APPROACH

The idea of contextual personalization, proposed and developed here, responds to the fact that human preferences are multiple, heterogeneous, changing, even contradictory, and should be understood in context with the user goals and tasks at hand [31]. Indeed, not all user preferences are relevant in all situations.

Context is a difficult notion to grasp and capture in a software system. In our approach, we focus our efforts on this major topic for content search and retrieval systems, by restricting it to the notion of semantic runtime context. The latter forms a part of general context, suitable for analysis in personalization and can be defined as the background themes under which user activities occur within a given unit of time. In this view, the problems to be addressed include how to represent the context, how to determine it at runtime, and how to use it to influence the activation of user preferences, *contextu-*

alize them and predict or take into account the drift of preferences over time (short and long term).

In our current solution to these problems, the runtime context is represented as (is approximated by) a set of weighted concepts from a domain ontology. This is built upon a personalization framework where user preferences are also considered to be concepts in the same domain [8]. Our approach to the contextual activation of preferences is then based on a computation of the semantic distance between each user preference and the set of concepts in the current context. This distance is assessed in terms of the number and length of the semantic paths linking preferences to context, across the semantic network defined by the ontology.

Ultimately, the perceived effect of contextualization is that user interests that are out of focus for a given context are disregarded, and only those that are in the semantic scope of the ongoing user activity (a sort of intersection between user preferences and runtime context) are considered for personalization. In practice, the inclusion or exclusion of preferences is not binary, but instead ranges on a continuum scale, where the contextual weight of a preference decreases monotonically with the semantic distance between the preference and the context.

Let us note that in the sequel the terms "preference" and "context" shall always refer to something implicit, as opposed to e.g. an explicit, literal user query. The user preferences handled by the system are assumed to be persistent, although we do not address here the issue of considering different durations or degrees of persistence. The dynamic acquisition, update and evolution of long-term user preferences (be they manual or automatic) are also considered as an external problem, which can be addressed independently from our present research.

A. Underlying Personalization Framework

The contextualization model presented here is grounded on an ontology-based personalization framework. Building on ontology-based semantic structures and semantic metadata, the personalization system builds and exploits an explicit awareness of (meta)information about the user, either directly provided by the user, or implicitly evidenced along the history of his/her actions.

The retrieval system assumes that the multimedia items in a retrieval space $\mathcal{D}$ are annotated with weighted semantic metadata which describe the meaning carried by the item content, in terms of a domain ontology $\mathcal{O}$. That is, each item $d \in \mathcal{D}$ is associated with a vector $M(d) \in [0,1]^{|\mathcal{O}|}$ of domain concept weights, where for each $x \in \mathcal{O}$, the weight $M_x(d)$ indicates the degree to which the concept x is important in the meaning of d.

The personalization system makes use of conceptual user profiles (as opposed to e.g. sets of preferred documents or keywords), where user preferences are represented as a vector of weights scaled between 0 and 1, corresponding to the intensity of user interest for each concept in the ontology. Comparing the metadata of items, and the preferred concepts in a user profile, the system predicts how the user may like an item, measured as a value in [0,1]. Based on this, contents (in a collection, a catalog section, a search result list, a video index, a structured multimedia object) are filtered and ranked in personalized ways. The reader is encouraged to find further details of this system in [8].

The ontology-based representation of user interests is richer, more precise, less ambiguous than a keyword-based or item-based model. It provides an adequate grounding for the representation of coarse to fine-grained user interests (e.g. interest for broad topics, such as football, sci-fi movies, or the NASDAQ stock market, vs. preference for individual items such as a sports team, an actor, a stock value), and can be a key enabler to deal with the subtleties of user preferences, such as their dynamic, context-dependent relevance.

An ontology provides further formal, computer-processable meaning on the concepts (who is coaching a team, an actor's filmography, financial data on a stock), and makes it available for the personalization system to take advantage of. Furthermore, current ontology standards, such as RDF [4] and OWL [22], support inference mechanisms that can be used in the system to further enhance personalization, so that, for instance, a user interested in animals (superclass of cat) is also recommended images showing cats. Inversely, a user interested in lizards, snakes, and chameleons can be inferred to be interested in reptiles with a certain confidence. Also, a user keen on Sicily can be assumed to like Palermo, through the transitive *locatedIn* relation.

B. Semantic Context for Personalization

Our model for context-based personalization can be formalized in an abstract way as follows, without any assumption on how preferences and context are represented. Let $\mathcal{U}$ be the set of all users, let $\mathcal{C}$ be the set of all contexts, and $\mathcal{P}$ the universe 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 each step in a session with the system, which we shall represent by a mapping $C : \mathcal{U} \times \mathbb{N} \to \mathcal{C}$, since we assume that the 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 $t \in \mathbb{N}$.

Definition 1. Let $\mathcal{C}$ be the set of all contexts, and let $\mathcal{P}$ be the set of all possible user preferences. 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 $c \in \mathcal{C}$, $p \models \Phi(p,c)$.

In this context the entailment $p \models q$ 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 preference was p.

Now we can particularize the above definition for a specific representation of preference and context. As explained in the previous section, in our model user preferences are represented by a set of weighted domain ontology concepts for which the user has an interest, where the intensity of the interest can range from 0 to 1.

Definition 2. Given a domain ontology $\mathcal{O}$, we define the *set of all preferences* over $\mathcal{O}$ as $\mathcal{P}_{\mathcal{O}} = [0,1]^{|\mathcal{O}|}$, 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.

Definition 3. Under the above definitions, we particularize $\models_{\mathcal{O}}$ as follows: given $p, q \in \mathcal{P}_{\mathcal{O}}$, $p \models_{\mathcal{O}} q \Leftrightarrow \forall x \in \mathcal{O}$, either $q_x \leq p_x$, or q_x can be deduced from p using consistent preference extension rules over $\mathcal{O}$. By extension rules we mean any formal procedure (e.g. logic, bayesian, statistic, or simply heuristic) that infers new preferences from an initial preference set, according to some stated theory or principle.

Now, our particular notion of context is that of the *semantic runtime context*, which we define as the background themes under which user activities occur within a given unit of time.

Definition 4. Given a domain ontology $\mathcal{O}$, we define the *set of all semantic runtime contexts* as $\mathcal{C}_{\mathcal{O}} = [0,1]^{|\mathcal{O}|}$.

With this definition, a context is represented as a vector of weights denoting the degree to which a concept is related to the current activities (tasks, goals, short term needs) of the user.

Note that although the definitions above will be particularized on a specific framework for personalized retrieval, we have not made any assumption so far on the type of application where the abstract model is to be implemented. Therefore, the proposed formalization is quite general. The model will be instantiated in the next sections, where we shall 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 will be given by $\Phi(P(u), C(u,t))$.

C. Building a Dynamic Retrieval Context

The model defined in the previous subsection is now particularized for content retrieval as follows. In the frame of a content retrieval system, we define the *semantic retrieval runtime user context* as the set of concepts that have been involved, directly or indirectly, in the interaction of a user u with the system during a retrieval session. Therefore, at each point t in time, we represent the retrieval context $C(u,t)$ as a vector in $[0,1]^{|\mathcal{O}|}$ of concept weights, where each $x \in \mathcal{O}$ is assigned a weight $C_x(u,t) \in [0,1]$. Time is measured by the number of user requests within a session. Since the fact that the context is relative to a user is clear, in the following we shall often omit this variable and use $C(t)$, or even C for short, as long as the meaning is clear.

In our approach, $C(t)$ is built as a cumulative combination of the concepts involved in successive user requests, in such a way that the importance of concepts fades away with time. This simulates the natural drift of user focus over time. Right after each user's request, a request vector $R(t) \in \mathcal{C}_{\mathcal{O}}$ is defined. This vector may be defined as the vector of concepts in the query, if the request consists of a query. In this case, the concepts can be extracted from a natural language or keyword-based query, using state of the art Information Extraction techniques [26]. If the request is of the type "view document", $R(t)$ can be defined by the topmost relevant concepts that annotate the document. If the request is a relevance feedback iteration step, $R(t)$ can be the average concept-vector corresponding to the set of documents marked as relevant by the user. Similar strategies can be defined to build concept vectors from browsing requests by topics and categories of documents or concepts, and other common content retrieval user action types.

Next, an initial context vector $C(t)$ is defined by combining the newly constructed request vector $R(t)$ with the context $C(t-1)$ computed in the previous step, where the context weights computed in step $t-1$ are automatically reduced by a decay factor $\xi \in [0,1]$. Consequently, at a given time t, we update $C_x(t)$ as:

$$C_x(t) = \xi \cdot C_x(t-1) + (1-\xi) \cdot R_x(t).$$

To the extent that $R(t)$ may contain concepts from search results selected by the user, this may seem similar to a relevance feedback strategy [6], [18]. However, here the context vector $C(t)$ is not used to reformulate the query, but to focus the preference vector, as shown next.

D. Contextual Preference Activation

Once a representation of the general user preferences and the live context are available, the selective activation of user preferences is based on finding semantic paths between preference and context concepts. The considered paths are made of semantic relations between concepts in the domain ontology, which form a semantic network. The shorter, stronger, and more numerous such connecting paths are, the more in context a preference will be considered. The semantic paths are explored by a form of Constraint Spreading Activation [10]. Our strategy consists of a semantic expansion of both user preferences and the context, during which the involved concepts are assigned preference weights and contextual weights, which decay as the expansion progresses farther away from the initial sets. This process can also be interpreted a sort of fuzzy semantic intersection between user preferences and the semantic runtime context, where the final computed weight of each concept represents the degree to which it belongs to each set (see Fig. 1).
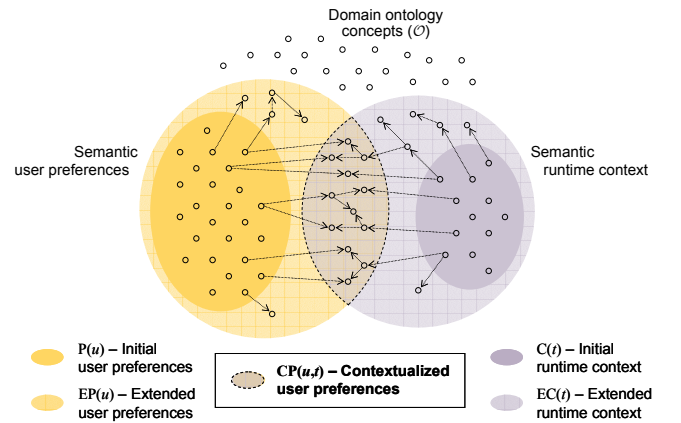

Fig. 1. Contextual activation of semantic user preferences.

For the propagation method each semantic relation r in the ontology is weighted by a propagation strength $w(r)$, which represents the likelihood that if we know that a concept x is in a certain context (or set of preferences), and $r(x,y)$ holds, (i.e. concepts x and y are related through r in the ontology), then y

should also be considered as part of this context. Based on these weights, our strategy spreads the initial context C (t) to a larger context vector EC(t) through the semantic network of semantic relations, where of course $\text{EC}_x(t) \geq \text{C}_x(t)$ for all $x \in \mathcal{O}$.

Let $\mathcal{R}$ be the set of all relations in $\mathcal{O}$, let $\overline{\mathcal{R}} = \mathcal{R} \cup \{r^{-1} \mid r \in \mathcal{R}\}$, and $w : \overline{\mathcal{R}} \rightarrow [0,1]$. The precise expression by which the extended context vector $EC(t)$ is computed is the following:

$$\text{EC}_y(t) = \begin{cases} \text{C}_y(t) \text{ if } \text{C}_y(t) > 0 \\ \Upsilon\left(\left\{\text{EC}_x(t) \cdot w(r) \cdot \text{pow}(x)\right\}_{x \in \mathcal{O}, r \in \overline{\mathcal{R}}, r(x,y)}\right) \text{ otherwise} \end{cases}$$

where $\text{pow}(x) \in [0,1]$ is a propagation power assigned to each concept x (by default, $\text{pow}(x) = 1$), allowing a finer control of the propagation through certain concepts (e.g. inhibition of propagation through very abstract concepts). Note that the top line of the expression above explicitly excludes the propagation between concepts in the input context (i.e. these remain unchanged after propagation). The Υ function is defined as follows. Given $X = \{x_i\}_{i=0}^n$, where $x_i \in [0,1]$,

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

This computation is based on the inclusion-exclusion principle applied to probability [33], where, put informally, $\text{EC}_y(t)$ would correspond to the probability that y is part of the context, and would be estimated in terms of the probability that other concepts y are in the context, where $w(r)$ would correspond to the conditional probability that y is in the context provided that x is in the context when $r(x,y)$ is known to be true.

After the context and preferences are expanded, only the preferred concepts with a context value different from zero (or above a threshold) shall count for personalization. This is done by computing a contextual preference vector CP, as defined by $\text{CP}_x = \text{EP}_x \cdot \text{EC}_x$ 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. $\text{CP}(u,t)$.

Note also that at this point we have achieved a contextual preference mapping Φ as defined in Section III.B, namely $\Phi(\text{P}(u),\text{C}(u,t)) = \text{CP}(u,t)$, where $\text{P}(u) \models \Phi(\text{P}(u),\text{C}(u,t))$, since $\text{CP}_x(u,t) > \text{P}_x(u,t)$ only when $\text{EP}_x(u)$ has been derived from $\text{P}(u)$ through the spreading mechanism, and $\text{CP}_x(u,t) < \text{EP}_x(u)$.

E. Personalized Retrieval in Context

Finally, given a multimedia document $d \in \mathcal{D}$ ($\mathcal{D}$ being the set of all documents in the retrieval space, as described in Section III.A), the predicted interest (to which we shall refer as *personal relevance measure*, prm) of the user u for d at a given instant t in a session is measured as a value in $[0,1]$ computed by:

$$\text{prm}(d,u,t) = \cos(\text{CP}(u, t-1), \text{M}(d)),$$

where $\text{M}(d) \in [0,1]^{|\mathcal{O}|}$ is the semantic metadata concept-vector of the document, as explained in Section III.A, whereby the simi-

larity between content descriptions and contextual preferences is measured as the cosinus of the angle formed by the corresponding vectors. In the context of a content retrieval system, where users retrieve contents by issuing explicit requests and queries, the prm measure is combined with query-dependent, user-neutral search result rank values, to produce the final, contextually personalized, rank score for the document:

$$\text{score}(d,q,u,t) = f(\text{prm}(d,u,t), \text{sim}(d,q)),$$

where the similarity measure $\text{sim}(d,q)$ stands for any ranking technique to rank documents with respect to a query or request. In general, the combination above can be used to introduce a personalized bias into any ranking technique that computes $\text{sim}(d,q)$, which could be image-based, ontology-based, relevance-feedback based, etc. The combination function f can be defined for instance as a linear combination $f(x,y) = \lambda \cdot \overline{x} + (1 - \lambda) \ \overline{y}$. The term λ is the personalization factor that shall determine the degree of personalization applied to the search result ranking, ranging from $\lambda=0$ producing no personalization at all, to $\lambda=1$, where the query is ignored and results are ranked only on the basis of global user interests. As a general rule, λ should decrease with the degree of uncertainty about user preferences, and increase with the degree of uncertainty in the query. The problem of how to set the value of λ dynamically is addressed by the authors in [8], where the reader is encouraged to find further details. $\overline{x}$ and $\overline{y}$ denote the normalization of the score values x and y, which is needed before the combination to ensure that they range on the same scale [13]. The final value $\text{score}(d,q,u,t)$ determines the position of each document d in the final ranking in the personalized search result presented to the user.

IV. A USE CASE

As an illustration of the application of the contextual personalization techniques, consider the following scenario: Clio's family and friends have set up a common repository where they upload and share their pictures and videos. Clio has not checked out the collection for quite a while, and she is willing to take a look at what images her relatives have brought from their last summer vacations.

Let us assume that the proposed framework has learned some of Clio's preferences over time, i.e. Clio's profile includes the weighted semantic interests for domain concepts of the ontology, shown in Table I, where *Tobby* is her brother's pet, and all the weights have been taken as 1.0 to simplify the example. This defines the P vector.

TABLE I
USER PREFERENCES

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

In our approach, these concepts are defined in a domain ontology containing other concepts and the relations between them, a subset of which is exemplified in Fig. 2.


Fig. 2. A subset of a domain ontology containing the concepts involved in the use case.

The propagation weight manually assigned to each semantic relation is shown in Table II. Weights were initially set by manually analyzing and checking the effect of propagation on a list of use cases for each relation, and was tuned empirically afterwards. Investigating methods for automatically learning the weights is an open research direction for our future work.

TABLE II
PROPAGATION WEIGHTS OF SEMANTIC RELATIONS

Relation r	$w(r)$	$w(r^{-1})$
contains	0.6	0.5
instanceOf	1.0	0.3
madeOf	0.7	0.6
similarTo	0.8	0.8
subclassOf	1.0	0.3

When Clio enters a query q_1 (the query-based search engine can be seen essentially as a black box for our technique), the personalization system adapts the result ranking to Clio's preferences by combining the query-based $sim(d,q_1)$ and the preference-based $prm(d,Clio)$ scores for each multimedia document d that matches the query, as described in Section III.E. At this point, the adaptation is not contextualized, since Clio has just started the search session, and the runtime context is still empty (i.e. at $t = 0$, $C(0) = \varnothing$). But now suppose that Clio selects and downloads an image and one video sequence shown in Fig. 3 from the search results.

Fig. 3. A picture (left) and a video sequence (right) selected by the user.

As a result, the system builds a runtime context out of the metadata of the selected documents, including the elements shown in Table III. This corresponds to the C vector (which for $t = 1$ is equal to $R(t)$), as defined in Section III.C.

TABLE III
CONTEXT VECTOR

C(Clio,1)	
Construction	1.0
Flower	1.0

Now, Clio wants to see some picture of her family members, and issues a new query q_2. The contextualization mechanism comes into place, as follows.

1. First, the context set is expanded through semantic relations from the initial context, adding two more weighted concepts, shown in bold in Table IV. This makes up the EC vector, following the notation used in Section III.D.

TABLE IV
EXPANDED CONTEXT

EC(Clio,1)	
Construction	1.0
City	**0.6**
Flower	1.0
Vegetation	**0.4**

2. Similarly, Clio's preferences are extended through semantic relations from her initial ones. The expanded preferences stored in the EP vector are shown in Table V, where the new concepts are in bold.

TABLE V
EXTENDED USER 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**		

3. The contextualized preferences are computed as described in section III.D, by multiplying the coordinates of the EC and the EP vectors, yielding the CP vector shown in Table VI (concepts with weight 0 are omitted).

TABLE VI
CONTEXTUALIZED USER PREFERENCES

CP(Clio,1)	
Construction	0.7
City	0.6
Flower	1.0
Vegetation	0.4

Comparing this to the initial preferences in Clio's profile, we can see that *Car*, *Sea* and *Toby* 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. Using the contextualized preferences above, a different personalized ranking is computed in response to the current user query q_2 based on the EC(Clio,1) vector, instead of the basic P(Clio) preference vector, as defined in Section III.E.

This example illustrates how our method can be used to contextualize the personalization in a query-based content search system, where the queries could be of any kind: visual ones, keyword-based, natural language queries. The approach could be similarly applied to other types of content access services, such as personalized browsing capabilities for multimedia repositories, automatic generation of a personalized slideshow, generation of personalized video summaries (where video frames and sequences would be treated as retrieval units), etc.

V. EXPERIMENTAL RESULTS

The contextualization techniques described in the previous sections have been implemented in an experimental prototype, and tested on a medium-scale corpus. Evaluating personalization is known to be a difficult and expensive task [27], [34]. On top of that, a formal evaluation of the contextualization techniques may require a significant amount of extra feedback from users in order to measure how much better a retrieval system can perform with the proposed techniques than without them. For this purpose, it would be necessary to compare the performance of retrieval a) without personalization, b) with simple personalization, and c) with contextual personalization. In this case, the standard evaluation measures from the IR field require the availability of manual content ratings with respect to a) query relevance, b) query relevance *and* general user preference (i.e. regardless of the task at hand), and c) query relevance *and* specific user preference (i.e. constrained to the context of his/her task). This requires building a testbed consisting of a search space corpus, a set of queries, and a set of hypothetic context situations, where users would be required to provide ratings to measure the accuracy of search results. The latter means considering sequences of user actions defined *a priori*, which makes it more difficult to get a realistic user assessment, since in principle the user would need to consider a large set of artificial, complex and demanding assumptions.

A. Experimental Setup

As an initial approach, yet allowing meaningful observations, we present here the results of our experimentation of the contextualization techniques, aiming to test the feasibility, soundness, and technical validity of the defined models and algorithms, including medium-sized scalability tests on a corpus of significant size. The corpus consists of 145,316 multimedia documents (445MB) from the CNN web site,[1] plus the KIM domain ontology and knowledge base (KB) [19], publicly available as part of the KIM Platform, developed by Ontotext Lab[2], with minor extensions. 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 intui-

---

[1] http://dmoz.org/News/Online_Archives/CNN.com
[2] http://www.ontotext.com/kim

tive basis, and tuned empirically afterwards by running a few trials.

The retrieval system used for this experiment is a semantic search engine developed by the authors [7], which did not implement itself any personalization capability. In order to extract precision and recall figures, we have rated the document / query / preference / context tuples manually. Since the contextualization techniques are applied in the course of a session, a sequence of steps needs to be defined in order to put them to work. Therefore we have defined a set of ten short use cases as part of the evaluation set up. As an example, one of such scenarios is explained next, along with the results obtained both in the individual experiment (see Fig. 4), and on average over the whole set (Fig. 5).

B. A Test Scenario

The sample scenario goes as follows. Alexander is fond of all kinds of luxurious and stylish articles. His 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. Alexander starts a search session with a query for news about *Daimler-Chrysler* and the different brands the company owns. *Daimler-Chrysler* owns luxury brands as *Mercedes* or *Maybach*, and other more ordinary ones like *Dodge* or *Setra* that are not of interest to Alexander.

Whereas the retrieval system does not rank the luxury brands any higher than the others, personalization reorders the results according to Alexander's preferences, showing first the documents related to *Daimler-Chrysler* and its higher-end brands *Mercedes* or *Maybach*, and pushing down other documents related to the lower-end brands of the company. In consequence, the personalized search performs better from the user's point of view. Since this is the first query of the session, no context exists yet, so user preferences are not filtered, and there is no contextualization performance to measure.

Now Alexander opens some documents in the search result, about the *Mercedes* brand and how *Daimler-Chrysler* is going to commercialize a new car model. He also opens a multimedia presentation about the new *Maybach 62* model. The context monitor extracts the concept of *Mercedes* from the annotated documents and images viewed by the user, along with the concept *Maybach*, since the selected content was mainly about these two brands. The context is expanded to new concepts such as *Daimler-Chrysler*, owner of *Mercedes* and *Maybach*, along with all its brands.

Next, Alexander makes a new query: "companies that trade on the New York Stock Exchange and have brands in the USA". The query results are re-ranked according to the contextualized preferences of Alexander. The documents that mention *Daimler-Chrysler* and *Mercedes* are pushed up in the result set, and the personal relevance increases also on the documents annotated with *Maybach*, the other *Daimler-Chrysler*'s favourite brand of the user. Alexander still encounters other companies and brands that trade in the New York stock exchange and match his preferences, like the *Sony Cor-*

poration, but these are not found semantically close to the brands in the context, and therefore get a lower ranking than other contents more in context with the previous user actions. This matches the real ongoing (implicit) user interests, which explains the improvement shown in Fig. 4.

C. Results and Discussion

The experiment described in the previous section is a clear example where personalization alone would not give better results, or would even perform worse than non-adaptive retrieval (see the drop of precision for recall between 0.1 and 0.3 in Fig. 4a), because irrelevant long term preferences (such as, in the example, the user's favourite luxury brands and companies which are not related to his current focus on the car industry context) would get in the way of the user. The experiment shows how our contextualization approach can avoid this effect and significantly enhance personalization by removing such out-of-context user interests and leave the ones that are indeed relevant in the ongoing course of action.
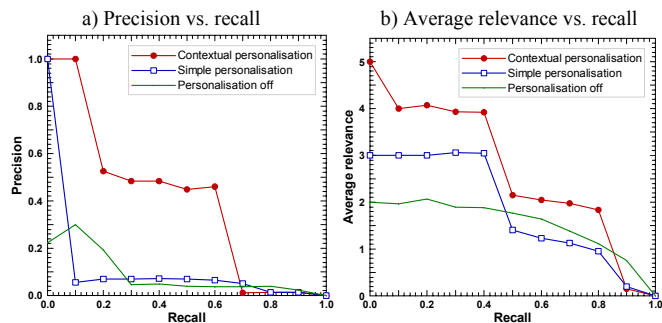

Fig. 4. Comparative performance of personalized search with and without contextualization, for the query "Companies that trade on the New York Stock Exchange and have a brand in the USA". The graphic a) shows the precision vs. recall curve, and b) shows the average relevance vs. recall.

The contextualization technique consistently results in better performance with respect to simple personalization, as can be observed in Fig. 5, which shows the average results over ten use cases. The cases where our technique performed worse were due to a lack of information in the KB, as a result of which the system did not find that certain user preferences were indeed related to the context. Another limitation of our approach is that it assumes that consecutive user queries tend to be related, which does not hold when sudden changes of user focus occur. However, not only the general improvements pay off on average, but the potential performance decay in such cases disappears after two or three queries, since the weight of contextual concepts decreases exponentially as the user keeps interacting with the system, as explained at the end of Section III.C. Nonetheless, as future work, it would be possible to enhance our approach by assessing the semantic distance between user requests, and clustering the context into cohesive subsets, leading to an even finer contextualization.
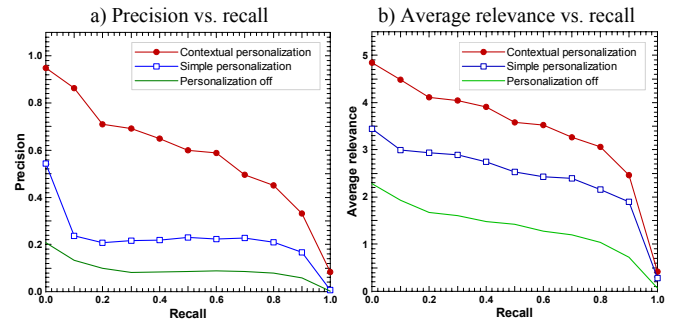

Fig. 5. Comparative performance of personalized search with and without contextualization, averaged over ten use cases.

In a way, our model of contextualized user preferences can be viewed as an approximation to short-term, live user interests, as opposed to the whole set of preferences, which would stand for the long-term ones. However, our model does not explicitly capture occasional short-term interests as such, unless they are persistently stored in the user profile. Since it is quite difficult to distinguish a casual, live user interest from a merely contextual concept, we include the former within the latter, in a way that in practice short-term user preferences can influence system responses. Still, if the implicit live interest is totally unrelated to the persistent preferences, its impact will be minimum or null. This is also an open research problem to be addressed in future work.

VI. CONCLUSION

Context is an increasingly common notion in IR. This is not surprising since it has been long acknowledged that the whole notion of relevance, at the core of IR, is strongly dependent on context – in fact it can hardly make sense out of it. Several authors in the IR field have explored approaches that are similar to ours in that they find indirect evidence of searcher interests by extracting implicit meanings in information objects manipulated by users in their retrieval tasks [3], [14], [16], [17], [21].

A first distinctive aspect in our approach is the use of semantic concepts, rather than plain terms (i.e. keywords), for the representation of these contextual meanings, and the exploitation of explicit ontology-based information attached to the concepts, available in a knowledge base. This extra, formal information allows one to determine the set of concepts that can be properly attributed to the context, in a more accurate and reliable way (by analyzing explicit semantic relations) than the statistical techniques used in previous proposals, which e.g. estimate term similarities by their statistic co-occurrence in a content corpus. Moreover, it allows the application of our techniques to multimedia corpora by means of semantic annotations which link the raw audiovisual content to the ontology-based conceptual space where user preferences and semantic context are modeled. Thus, our proposal can reap the benefits from automatic content analysis research [24].

Other than this, our approach is novel in that it combines the implicit context meanings collected at runtime, with a persistent, more general representation of user interests, learned by the system over a period of time or provided manually by the user, prior to a search session. The benefit is twofold: the 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 senseful by relating them to user preferences. This does not completely remove the uncertainty inherent to the prediction of implicit user interests involved in any approach to personalization, but it can significantly reduce inaccuracies in a considerable number of cases.

REFERENCES

[1] Abowd, G. D., Dey, A. K., Brown, P. J., Davies, N., Smith, M., and Steggles, P., "Towards a Better Understanding of Context and Context-Awareness," in *Proc. of the 1st Int. Symp. on Handheld and Ubiquitous Computing*, Springer Verlag LNCS vol. 1707, 1999, pp. 304-307.

[2] Akrivas, G., Wallace, M., Andreou, G., Stamou, G., and Kollias, S., "Context-Sensitive Semantic Query Expansion," in *Proc. of the IEEE International Conference on Artificial Intelligence Systems*, Divnomorskoe, Russia, 2002, pp. 109-104.

[3] Bharat, K., "SearchPad: Explicit capture of search context to support web search," in *Proc. of the 9th International World Wide Web Conf.*, Amsterdam, 2000, pp. 493-501.

[4] Brickley, D., and Guha, R. V., "RDF Vocabulary Description Language 1.0: RDF Schema," *W3C Recommendation* 10 February 2004, http://www.w3.org/TR/rdf-schema.

[5] Brown, P. J., Bovey, J., and Chen, X., "Context-Aware Applications: From the Laboratory to the Marketplace," *IEEE Personal Communications* vol. 4, no. 5, 1997, pp. 58-64.

[6] Campbell, I., and Van Rijsbergen, C. J., "The ostensive model of developing information needs," in *Proc. of the 2nd International Conference on Conceptions of Library and Information Science*, Copenhagen, Denmark: Royal School of Librarianship, 1996, pp. 251-268.

[7] Castells, P., Fernández, M., and Vallet, D., "An Adaptation of the Vector-Space Model for Ontology-Based Information Retrieval," *IEEE Transactions on Knowledge and Data Engineering* vol. 19, no.2, *special issue on Knowledge and Data Engineering in the Semantic Web Era*, 2007, pp. 261-272.

[8] Castells, P., Fernández, M., Vallet, D., Mylonas, Ph., and Avrithis, Y., "Self-Tuning Personalized Information Retrieval in an Ontology-Based Framework," in *Proc. of the 1st Int. Workshop on Web Semantics*, Springer Verlag LNCS vol. 3762, 2005, pp. 977-986.

[9] Chen, P. M., and Kuo, F. C., "An information retrieval system based on a user profile", *Journal of Systems and Software* vol. 54, no. 1, 2000, pp. 3-8.

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

[11] Della Mea, V., and Mizzaro, S., "Measuring retrieval effectiveness: a new proposal and a first experimental validation," *Journal of the American Society for Information Science and Technology* vol. 55, no. 6, 2004, pp. 530-543.

[12] Edmonds, B., "The Pragmatic Roots of Context", in *Proc. of the 2nd International and Interdisciplinary Conference on Modeling and Using Context*, Springer Verlag LNAI vol. 1688, 1999, 119-132.

[13] Fernández, M., Vallet, D., and Castells, P., "Probabilistic Score Normalization for Rank Aggregation," in *Proc. of the 28th European Conf. on Inf. Retrieval*, LNCS vol. 3936. Berlin Heidelberg New York: Springer Verlag, 2006, pp. 553-556.

[14] Finkelstein, L., Gabrilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolfman, G., and Ruppin, E., "Placing Search in Context: The Concept Revisited," *ACM Transaction on Information Systems* vol. 20 no. 1, 2002, pp. 116-131.

[15] Gauch, S., Chaffee, J., and Pretschner, A., "Ontology-based personalized search and browsing," *Web Intelligence and Agent Systems* vol. 1 no. 3-4, 2003, pp.219-234.

[16] Haveliwala, T. H., "Topic-Sensitive PageRank: A Context-Sensitive Ranking Algorithm for Web Search," *IEEE Trans. on Knowledge and Data Engineering* vol. 15 no. 4, 2003, pp. 784-796.

[17] Jones, G. J. F., and Brown, P. J., "The Role of Context in Information Retrieva,l" in *Proc. of the SIGIR Information Retrieval in Context Workshop,* Sheffield, UK: ACM Press, 2004, pp. 20-22.

[18] Kelly D., Teevan, J., "Implicit feedback for inferring user preference," *SIGIR Forum* vol. 32 no. 2, 2003, pp. 18-28.

[19] Kiryakov, A., Popov, B., Terziev, I., Manov, D., and Ognyanoff, D., "Semantic Annotation, Indexing, and Retrieval," *Journal of Web Semantics* vol. 2 no. 1, 2004, pp. 49-79.

[20] Kobsa, A., "Generic User Modelling Systems," *User Modeling and User-Adapted Interaction* vol. 11, 2001, pp. 49-63.

[21] Lawrence, S., "Context in Web Search," *IEEE Data Engineering Bulletin* vol. 23 no. 3, Sep. 2000, pp. 25-32.

[22] McGuinness, D. L., and van Harmelen, F., "OWL Web Ontology Language Overview," *W3C Recommendation* 10 February 2004, http://www.w3.org/TR/owl-features.

[23] Mylonas, Ph., and Avrithis, Y., "Context modeling for multimedia analysis and use," in *Proc. of the 5th International and Interdisciplinary Conference on Modeling and Using Context*, Paris, France, 2005.

[24] Papadopoulos, G. Th., Mylonas, Ph., Mezaris, V., Avrithis, Y., Kompatsiaris, I., "Knowledge-Assisted Image Analysis Based on Context and Spatial Optimization," *International Journal on Semantic Web and Information Systems* vol. 2 no. 3, 2006, pp. 17-36.

[25] Pascoe, J., "The Stick-e note architecture: Extending the interface beyond the user," in *Proc. of the International ACM Conference on Intelligent User Interfaces*, Orlando, FL: ACM Press, 1997, pp. 261-264.

[26] Popov, B., Kiryakov, A., Ognyanoff, D., Manov, D., and Kirilov, A., "KIM – A Semantic Platform for Information Extraction and Retrieval," *Journal of Natural Language Engineering* vol. 10 no. 3-4, 2004, pp. 375-392.

[27] Rajagopalan, B., and Deshmukh, A., "Evaluation of Online Personalization Systems: A Survey of Evaluation Schemes and A Knowledge-Based Approach," *Journal of Electronic Commerce Research* vol. 6 no. 2, 2005, pp. 112-122.

[28] Rocchio, J., "Relevance feedback information retrieval," in Salton, G., Ed., *The SMART Retrieval System: Experiments in Automatic Document Processing.* Englewood Cliffs, NJ: Prentice-Hall, 1971, pp. 313-323.

[29] Schilit, B., Adams, N., and Want, R., "Context-Aware Computing Applications," in *Proc. of the IEEE Workshop on Mobile Computing Systems and Applications*, Santa Cruz, CA, 1994, pp. 85-90.

[30] Staab, S., and Studer, R., Eds., "Handbook on Ontologies," *International Handbooks on Information Systems*, Berlin Heidelberg New York: Springer-Verlag, 2004.

[31] Vallet, D., Fernández, M., Castells, P., Mylonas, Ph., and Avrithis, Y., "Personalized Information Retrieval in Context," in *Proc. of the 3rd International Workshop on Modeling and Retrieval of Context at the 21st National Conference on Artificial Intelligence*, Boston, MA, 2006.

[32] White, R. W., Jose, J. M., and Ruthven, I., "An implicit feedback approach for interactive information retrieval," *Information Processing and Management* vol. 42 no. 1, 2006, pp. 166-190.

[33] Whitworth, W. A., *Choice and Chance, with one thousand exercises*, New York: Hafner Pub. Co., 1965.

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