

M

MULTIMEDIA PERSONALIZATION

Phivos Mylonas and Yannis Avrithis

National Technical University of Athens, Athens, Greece

Definition: *Multimedia personalization is tailoring digital audiovisual content to its users, based on personal details or characteristics they provide. It allows a content provider to adapt specific multimedia content in accordance with individual standards, tastes and preferences. For example, movies can be personalized based on genre, cast, director, etc.*

Introduction

During the last decade, the exceptional growth of ubiquitous communication technologies, side by side to the breathtaking increase of available digital multimedia content, has allowed access to personalized multimedia content *anytime, anywhere*. This growth comes with an increasing heterogeneity of client devices, as well as user preferences. The client devices range from traditional Personal Computers (PCs) to enhanced digital camera mobile phones and Personal Digital Assistants (PDAs) and all of them facilitate quite different device profiles in terms of resolution, computational power and memory. Therefore, the heterogeneity of clients and their ubiquitous connections pose new challenges in the delivery and presentation of rich personalized multimedia content to the end-users, according to their context, their specific background and interest, as well as their profiles and preferences.

The above phenomenon is strongly related to the fact that the cost of multimedia storage and wide area communication services has decreased, while at the same time their capacity increased dramatically. This observation, along with the increasing penetration of rich multimedia content applications, has made digital storage, annotation and access of multimedia information a mature and viable choice for content providers and individuals. Numerous multimedia collections have been made accessible, depending upon copyright, policy and security decisions, over the Internet in a cost-, time- and place-efficient fashion. However, one of the main problems of traditional content archiving and sharing has been inherited to its modern digital descendants and became even worse: people (i.e. the ultimate content users) are not interested in the entire set of available content, but quite on the contrary, they like to search and retrieve only specific portions of it, according to their personalized preferences and needs. Nowadays, they also get confused and overloaded by the vast amounts of multimedia content available, thus the need for multimedia personalization is now more evident than ever.

Dynamically created personalized content may include different modalities. Typically, *text-centric* personalized content may be identified on the Web and adaptive hypermedia systems [12]. On-demand generation and publication of personalized *multimedia* content is introduced with research efforts, like the personalized album MyPhotos [13] and services, like Flickr [14] or Youtube [22]. In the area of dynamically generating personalized multimedia presentations, [15] and [16] form indicative research examples, whereas SlideShare [17] introduced webinars and collaborative software to the broad audience. Recently, social networking services and mashups introduced a new era in the personalization process, both in personal (e.g. Facebook [18], Myspace [19], Flixster [25]) and professional (e.g. LinkedIn [20], Plaxo [21]) user space, by utilizing user generated tagging information and recommendations. In all cases, personalization seeks to improve the subjective performance of retrieval as perceived by individual users ([34], [35], [36], [37], [38]). This article focuses on the role of contextualized personalization in multimedia information retrieval (IR), and more specifically, in its smooth integration into the personalization of multimedia content retrieval.

In principle, personalized multimedia content access aims at enhancing the IR process by complementing explicit user requests with implicit user preferences, to better meet individual user needs [23]. Personalization is being currently envisioned as a major research trend to relieve the well-known *information overload* problem [24]. Typically, multimedia IR systems list relevant items only if users ask for them. Efficient multimedia personalization systems, on the other hand, should be able to predict the needs of users and recommend items, even though the user does not specifically request them. Thus, from the research point of view, multimedia personalization is typically separated into two main categories, namely *extraction of user preferences and profiles* and *personalized retrieval and ranking of the results*, to be discussed in the following.

The role of user profiles

It is a fact that uncertainty dominates the process of personalized information retrieval [1], as a limited set of terms cannot fully describe the user's wish. The role of multimedia personalization is to reduce this uncertainty, by using more information about the user's wishes than just the local interest. The contribution of user profiles in understanding the effect inherent in information retrieval, when two distinct users presenting identical queries obtain different subsets of retrieved documents and to different degrees, is crucial. The user profile is generated through the constant monitoring of the user's interaction, which contains less uncertainty because of the nature of his/her actions, as long as the monitoring period is sufficient and representative of the user's preferences. Therefore, a user profile, which contains valuable information concerning the user's global interest, i.e. information concerning the user's preferences over a long period of time, may be used whenever the query, i.e. the user's local preference or in other words the scope of his/her current interaction, provides insufficient information about the user and his/her local interest.

In order to process the user profile using the stored knowledge, the representation of the former needs to be compatible with it. From the semantic perspective, the best way to model the underlying knowledge is to utilize the notion of ontologies [2]. An ontology is "*an explicit specification of a conceptualization*" and it consists by several components of which the most important ones are concepts and relations. Such a (semantic) ontology usually contains semantic concepts and semantic relationships between these concepts,

whereas in the “fuzzified” case [3], degrees of confidence are associated to the corresponding relations between any two concepts, as well. We may define user preferences on the same set of concepts, or in other words, preferences are considered to be concepts, as well. When the user poses a query that is in fact related to one of his/her preferences, that preference may be used to facilitate the interpretation of the query, as well as the ranking of the selected multimedia documents. However, usage of preferences that are unrelated to the query may only be viewed as addition of noise, as any proximity between selected multimedia documents and these preferences is coincidental in the given context. Thus, in addition to positive preferences, special care must be taken for the representation and separate store of negative preferences, so that they are processed separately.

Usage history and user actions

In the process of identifying both kinds of user preferences, we start from the set of multimedia documents available in each user's usage history. Any user profile implementation, such as the one introduced in [4], may receive this usage history as input and produce the corresponding set of user preferences as output. In order to achieve this, the process needs to access the utilized knowledge (e.g. in terms of a semantic ontology). The set of multimedia documents available in the usage history is constructed as the result of the application of all user action types, during the user's interaction with the retrieval system. These actions characterize the user and express his/her personal view of the search space content. These actions are directly associated to user requests or queries and therefore we shall use the term *query* in the following. Typical multimedia content retrieval user action types that a user may pose as queries are:

- *keyword-based queries*: keywords may be extracted from a natural language or a keyword-based encountered query and are mapped to concepts in the annotation of multimedia documents, utilizing state-of-the-art information extraction techniques (e.g. [5]).
- *view multimedia document queries*: concepts are directly encountered in the annotation of a multimedia document and are usually retrieved with the help of a semantic index [6].
- *relevance feedback queries*: relevance feedback queries satisfy users' relevance feedback requests and consist mainly of two parts, namely positive and negative relevance feedback requests. Positive requests correspond to the annotation of the set of multimedia documents marked as relevant by the user, whereas negative ones correspond to the annotation of the set of multimedia documents marked as non relevant [26].
- *browsing content queries*: browsing queries are defined as a set of topics requested for browsing by the user, according to one specific browsing topic or category of documents or concepts.

Typically, the user's usage history comprises of a combination of all types of actions, provided that a user is able to perform any type of action at a given time. An association between the related history documents and concepts exists through the utilization of a semantic index [6], which is a priori constructed during analysis of either the raw content, or the associated textual annotation.

User preferences extraction

The formal definition of user preferences as a *fuzzy* set of concepts, allows participation of a single concept in multiple preferences and to different degrees¹. As already stated, the history of the user is typically represented as a (fuzzy) set on the set of concepts that are related to it and consists of both positive and negative parts. Preferences are mined by using both of these parts as input and by applying clustering algorithms on them. An ad-hoc methodology that further optimizes the above steps is the use of contextual information. Utilizing the notion of context in the process, we are able to extract two distinct sets of positive and negative user preferences as output and combine them in a meaningful way to obtain the final preferences. 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 ([27], [28], [29], [30], [31]). 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 [32].

Most clustering methods found in the literature belong to either of two general categories, partitioning or hierarchical [7]. Hierarchical methods do not require the number of clusters as input, in contrast to their partitioning counterparts. Since the number of preferences that may be encountered in a multimedia document is not known beforehand, the latter are inapplicable [8]. The same applies to the use of a supervised clustering method which allows one concept to belong to two or more clusters, such as fuzzy c-means [9]; the algorithm requires the number of concept clusters as input, i.e. it uses a hard termination criterion on the amount of clusters and, thus, can not be adopted for efficient user preferences extraction. In the following we shall outline an innovative hybrid approach, based on the fuzzification of an agglomerative² hierarchical clustering algorithm (Figure 1).

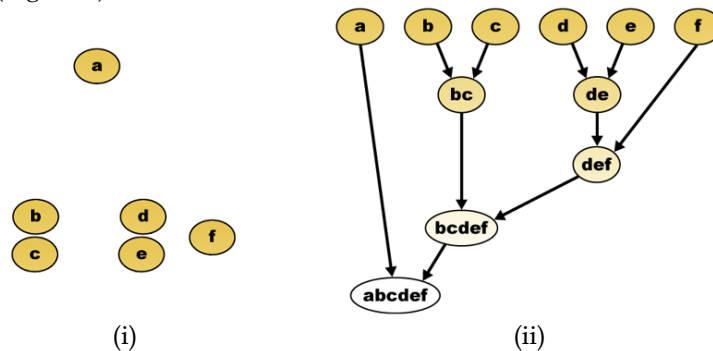


Figure 1. Hierarchical clustering of concepts used in multimedia personalization; (i) concepts to be clustered, (ii) clustering steps.

¹ Given a universe V , a crisp set S of concepts on V is described by a membership function $\mu_S: V \rightarrow \{0,1\}$. The crisp set S is defined as $S = \{s_i\}$, $i=1, \dots, N$. A *fuzzy* set F on S is described by a membership function $\mu_F: S \rightarrow [0,1]$.

² Hierarchical clustering methods are divided into agglomerative and divisive. The former are more widely studied and applied, as well as more robust and therefore are followed herein.

The application of such a clustering algorithm results into a crisp set of clusters detected in a user's usage history; each cluster is a crisp set of concepts. However, this alone is not sufficient, as it is rather obvious that such an approach is required to support multimedia documents belonging to multiple distinct preferences by different degrees. At the same time, the robustness and efficiency of the hierarchical clustering approach should also be retained, thus, without any loss of functionality or increase of computational cost, the crisp clusters are replaced by fuzzy normalized clusters. A fuzzy classifier [33] from the set of crisp to the set of fuzzy clusters of concepts is constructed. The fuzzy set of preferences associated to each fuzzy cluster is then obtained, by exploiting its context and cardinality information. Then, by aggregating the process to the entire set of fuzzy clusters, the fuzzy set of preferences related to the initial set of multimedia documents in the user's usage history is identified, after limiting it according to the predefined set of all possible user preferences.

The rest of this section provides analytical details on the initial concept clustering process, the cluster fuzzification, as well as the final user preference extraction. This threefold model can be formalized in an abstract way as a function $Y = G(X)$ without any assumption on how the input or output of the function may be represented and instantiated. The function takes a fuzzy set X as input and provides a different fuzzy set Y as its output. In this context, we may particularize the above statement for the specific case of positive user preferences P^+ and positive usage history H^+ ; function G can be utilized to obtain P^+ from H^+ , as: $P^+ = G(H^+)$. The described approach may then be decomposed into the following four general steps:

1. Perform a crisp clustering of concepts H^+ in order to determine the count of distinct positive preferences P^+ that a history document is related to
2. Construct a fuzzy classifier that measures the degree of correlation of a concept s with each crisp cluster.
3. Consider the context and cluster cardinality of the resultant fuzzy clusters and mathematically adjust their computed values so as to match their semantically anticipated counterparts.
4. Identify the positive user preferences P^+ that are related to each cluster, according to the *a priori* known set of all possible user preferences, in order to acquire an overall result.

The same applies in the case of the application of function G to H^- , in order to obtain P^- as: $P^- = G(H^-)$. As already stated, the final set of preferences P that correspond to the user's history is the set of positive P^+ meaningfully combined with the set of negative preferences P^- . Using the sum notation for fuzzy sets [8], this may be represented as: $P = \sum_{s \in S} s / \max(0, P^+(s) - P^-(s))$, where $P(s) = \max(0, P^+(s) - P^-(s))$ denotes

the final preference membership degree for each concept s . Finally, the set of preferences that correspond to the set of history documents associated to the user queries is the set of preferences that belong to any of the detected clusters of concepts that index the given multimedia documents. An illustrative mock-up example is given in Figure 2.




Figure 2. Fuzzy preferences extraction mock-up; concepts/topics are denoted by blue bubbles.

As observed in the figure, the above left set corresponds to the set of preferences related to the first cluster and the above right set is the set of preferences related to the second cluster. The set of preferences that belong to any of the two clusters is given by their union $prefs_I \cup prefs_{II}$, i.e. the set of all three top concepts. In order to validate the results of the above fuzzy classification, we have to assure that the set of topics that correspond to the set of documents H^+ are derived from the predefined set of all possible actual user preferences. Thus, we compute the intersection of both sets, resulting into the limitation of the final set of user preferences to the two shaded topics indicated in the Figure.

Personalized retrieval and ranking

Now, given a multimedia document of the retrieval space, the predicted interest (to which we shall refer as *personal ranking measure*, r_p) of the user for it at a given time instant in a session is measured as a value in the interval $[0,1]$, based on the previously extracted preferences and computed by:

$$r_p(doc, time) = \cos(S_{doc}, P_{time-1}) = \frac{|S_{doc} \cap P_{time-1}|}{\sqrt{|S_{doc}| |P_{time-1}|}}$$

where S_{doc} is the fuzzy set of concepts associated to the document, P_{time-1} the fuzzy set of (contextualized) preferences obtained from the previous subsection and min is the fuzzy intersection of the two. In the context of a multimedia content retrieval system, where users retrieve contents by issuing explicit requests and queries, the above measure is combined with query-dependent, user-neutral search result rank values, to produce the final, contextually personalized, rank score for the multimedia document. The *final personalized rank score* $r(doc, time)$ for the document is then given by the function:

$$r(doc, time) = f(r_p(doc, time), r_s(doc, time))$$

The *similarity measure*, $r_s(doc, time)$, which stands for any ranking technique to rank a document with respect to a query or request at a given time, is computed according to

the given possible user queries. For instance, in the case of keyword-based queries Q_k we have $r_s(doc, time) = \cos(S_{doc}, Q_k)$, or in the case of topic browsing, the degree to which the document is classified to topic Q_b is given by $r_s(doc, time) = G_{doc}(z)$, where $z = Q_b$ is the specific topic and $G_{doc} = G(S_{doc})$ is the topic classification output of the topic classification process. Both processes of user preferences extraction and topic classification implement the same algorithm and can be defined in terms of the same function G :

- $G(S_{doc}) = G_{doc}$ provides the fuzzy set of all topics associated to the specific document,
- z corresponds to a specific topic element of this fuzzy set, and
- $G_{doc}(z)$ denotes the degree to which z belongs to G_{doc} .

Multimedia documents are ranked according to their similarity to the predefined topic of search, whereas in the case of a single view document query, the requested document is simply presented to the user.

SEMANTIC AND METADATA SEARCH - AssembledResponse			
The responses from each archive have been assembled into the following list of retrieved documents, with degree of relevance:			
Id	Title	SourceArchive	Score
35	Flugzeugkatastrophe	FAA	1.0
21	Julius Raab tot	FAA	1.0
52	Die Vietnamkrise	FAA	1.0
1199	Επικυρώσεις Αυγούστου 1974	ERT	1.0
1514	Περικσότιο	ERT	1.0
AVQ-A-004129-0038	Archeological excavations in Rome	Alinari	1.0
FCC-F-021960-0000	Exodus of the Belgian population	Alinari	1.0

Pages:	≤ Previous	ArchivesResponse	AssembledResponse	PresentationResponse	ClassificationResponse	Next >=
--------	------------	-------------------------	--------------------------	----------------------	------------------------	---------

Figure 3. The retrieval response has been assembled into one list of retrieved multimedia documents; a fuzzy degree of relevance is imposed.

In general, the *final personalized rank score* can be used to introduce a personalized bias into any ranking technique that computes the *similarity measure*, which could be image-based, ontology-based, relevance-feedback based, etc. The combination function f can be defined for instance as a linear combination $f(x, y) = \lambda \cdot \bar{x} + (1 - \lambda) \cdot \bar{y}$. The term λ is the personalisation factor that shall determine the degree of personalisation applied to the search result ranking, ranging from $\lambda = 0$ producing no personalisation 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 dynamically is addressed for instance in [11], where the reader is encouraged to find further details. \bar{x} and \bar{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. The final value of the *personalized rank score* determines the position of each multimedia document in the final ranking in the personalized search result presented to the user. Typically the user has the possibility to see details of the retrieval process, e.g. the semantic concepts that matched her/his query, the simple ranked list of her/his results (Figure 3), as well as the final re-ranked list of the multimedia documents after the impact of the user's personalization settings, as shown in Figure 4.

SEMANTIC AND METADATA SEARCH - ClassificationResponse			
The search results have been additionally sorted according to your semantic preferences (semantic interests and thematic category preferences). The updated list of retrieved documents is the following:			
Id	Title	SourceArchive	Score
1199	Επικρατοίητες Αυγούστου 1974	ERT	0.9
21	Julius Raab tot	FAA	0.85
35	Flugzeugkatastrophe	FAA	0.8
AVQ-A-004129-0038	Archeological excavations in Rome	Alinari	0.7
1514	Περικσόμιο	ERT	0.65
52	Die Vietnamkrise	FAA	0.5
FCC-F-021960-0000	Exodus of the Belgian population	Alinari	0.45

Pages:	<< Previous	ArchivesResponse	AssembledResponse	PresentationResponse	ClassificationResponse
--------	-------------	------------------	-------------------	----------------------	------------------------

Figure 4. The search results have been sorted according to the user's semantic preferences (semantic interests and preferences).

References

1. C. H. Chang and C. C. Hsu, "Integrating query expansion and conceptual relevance feedback for personalized Web information retrieval," *Computer Networks and ISDN Systems*, vol. 30, pp. 621-623, 1998.
2. T. R. Gruber, "A translation approach to portable ontologies," *Knowledge Acquisition*, 5(2), pp. 199-220, 1993.
3. Ph. Mylonas, D. Vallet, M. Fernández, P. Castells and Y. Avrithis, "Ontology-based Personalization for Multimedia Content," 3rd European Semantic Web Conference (ESWC) - Semantic Web Personalization Workshop, Budva, Montenegro, 11-14 June 2006.
4. Ph. Mylonas, D. Vallet, P. Castells, M. Fernandez and Y. Avrithis, "Personalized information retrieval based on context and ontological knowledge," *Knowledge Engineering Review*, Cambridge University Press, Volume 23, Issue 1, 2008.
5. B. Popov, A. Kiryakov, D. Ognyanoff, D. Manov and A. Kirilov, "KIM - A Semantic Platform for Information Extraction and Retrieval," *Journal of Natural Language Engineering*, 10(3-4), pp. 375-392, 2004.
6. D. Vallet, Ph. Mylonas, M. A. Corella, J. M. Fuentes, P. Castells and Y. Avrithis, "A Semantically-Enhanced Personalization Framework for Knowledge-Driven Media Services," *IADIS International Conference on WWW / Internet (ICWI '05)*, Lisbon, Portugal, October 19-22, 2005.
7. S. Theodoridis and K. Koutroumbas, "Pattern Recognition," Academic Press, 1998.
8. S. Miyamoto, "Fuzzy Sets in Information Retrieval and Cluster Analysis," Kluwer Academic Publishers, Dordrecht Boston London, 1990.
9. M. Benkhalifa, A. Bensaid and A. Mouradi, "Text categorization using the semi-supervised fuzzy c-means algorithm," 18th International Conference of the North American Fuzzy Information Processing Society - NAFIPS, pp. 561-568. 1999.
10. L. Egghe and C. Michel, "Construction of weak and strong similarity measures for ordered sets of documents using fuzzy set techniques," *Information Processing and Management* vol. 39, no. 5, September 2003, pp. 771-807.
11. P. Castells, M. Fernandez, D. Vallet, Ph. Mylonas and Y. Avrithis, "Self-Tuning Personalized Information Retrieval in an Ontology-Based Framework," 1st Int. Workshop on Web Semantics, Springer Verlag LNCS vol. 3762, 2005, pp. 977-986.

12. P. Brusilovsky, A. Kobsa, and J. Vassileva, "Adaptive Hypertext and Hypermedia," Kluwer Acad. Publ., Dordrecht, 1998.
13. Y. Sun, H. Zhang, L. Zhan, and M. Li, "MyPhotos - A System for Home Photo Management and Processing," ACM Multimedia Conf. (ACMMM 2002), Juan-les-Pins, France, Dec. 1-6 2002.
14. Flickr, <http://www.flickr.com/>
15. Joost Geurts, Stefano Bocconi, Jacco van Ossenbruggen, and Lynda Hardman, "Towards Ontology-driven Discourse: From Semantic Graphs to Multimedia Presentations," Second International Semantic Web Conference (ISWC2003), October 20-23, 2003, Sanibel Island, Florida, USA, <http://homepages.cwi.nl/~media/cuyppers/>
16. M. J. F. Bes and F. Khantache, "A Generic Architecture for Automated Construction of Multimedia Presentations," Intl. Conf. on Multimedia Modeling (MMM 2001), Amsterdam, The Netherlands, Nov. 5-7 2001.
17. SlideShare, <http://www.slideshare.net/>
18. FaceBook, <http://www.facebook.com/>
19. MySpace, <http://www.myspace.com/>
20. LinkedIn, <http://www.linkedin.com/>
21. Plaxo, <http://www.plaxo.com/>
22. Youtube, <http://www.youtube.com/>
23. Kobsa, A., "Generic User Modelling Systems," User Modeling and User-Adapted Interaction vol. 11, 2001, pp. 49-63.
24. Chen, P. M., and Kuo, F. C., "An information retrieval system based on a user profile," Journal of Systems and Software vol. 54, no. 3-8, 2000.
25. Flixster, <http://www.flixster.com/>
26. 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.
27. Bharat, K., "SearchPad: Explicit capture of search context to support web search," in Proc. of the 9th International World Wide Web Conf., Amsterdam, 2000.
28. Lawrence, S., "Context in Web Search," IEEE Data Engineering Bulletin vol. 23 no. 3, 2000, pp. 25-32.
29. Finkelstein, L., Gabrilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolf-man, G., and Ruppin, E., "Placing Search in Context: The Concept Re-visited," ACM Transaction on Information Systems vol. 20 no. 1, 2002, pp. 116-131.
30. 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.
31. Jones, G. J. F., and Brown, P. J., "The Role of Context in Information Retrieval" in Proc. of the SIGIR Information Retrieval in Context Work-shop, Sheffield, UK: ACM Press, 2004
32. 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.
33. Kuncheva L.I. , "Fuzzy Classifier Design," Springer-Verlag, Heidelberg, May 2000.

34. Jeh, G. and Widom, J., "Scaling personalized web search," in Proceedings of the 12th International World Wide Web Conference (WWW 2003), Budapest, Hungary, pp. 271-279, 2003.
35. Gauch, S., Chaffee, J. and Pretschner, A., "Ontology-based personalized search and browsing," *Web Intelligence and Agent Systems* 1(3-4), 219-234, 2004.
36. Liu, F., Yu, C. and Meng, W., "Personalized web search for improving retrieval effectiveness," *IEEE Transactions on Knowledge and Data Engineering* 16(1), 28-40, 2004.
37. Micarelli, A. and Sciarrone, F., "Anatomy and empirical evaluation of an adaptive web-based information filtering system," *User Modelling and User-Adapted Interaction* 14(2-3), 159-200, 2004.
38. Castells, P., Fernandez, M., Vallet, D., Mylonas, Ph. and Avrithis, Y., "Self-tuning personalized information retrieval in an ontology-based framework," in Proceedings of the 1st International Workshop on Web Semantics, Agia Napa, Cyprus, Springer Verlag LNCS vol. 3762, pp. 977-986, 2005.