

# Personalized Information Retrieval in Context

David Vallet,<sup>\*</sup> Miriam Fernández,<sup>\*</sup> Pablo Castells,<sup>\*</sup> Phivos Mylonas,<sup>†</sup> and Yannis Avrithis<sup>†</sup>

<sup>\*</sup>Universidad Autónoma de Madrid, Escuela Politécnica Superior  
Campus de Cantoblanco, 28049 Madrid, Spain  
{david.vallet, miriam.fernandez, pablo.castells}@uam.es

<sup>†</sup>School of Electrical and Computer Engineering, National Technical University of Athens  
9, Iroon Polytechniou str., 15773 Zographou, Athens, Greece  
{fnylonas, iavr}@image.ntua.gr

## 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 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.

## 1. Introduction

The size and the pace of growth of the world-wide body of available information in (text and a/v) digital format constitute a permanent challenge for content retrieval technologies. People have instant access to unprecedented inventories of content world-wide, just a few clicks away 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 global organization, with decentralized content provision, dynamic networks, etc., query-based and browsing technologies often find their limits.

Personalized 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 (Gauch et al 2003). The combination of long-term and short-term user interests that takes place in this 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 differ-

ential aspects of individual users while avoiding to distract them away from their current specific goals.

Reliability is indeed a well-known concern in the area of user modeling and personalization technologies. 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, 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. This paper is concerned with exploiting semantic, ontology-based contextual information, specifically aimed towards its use in personalization for information retrieval. The goal of the research presented here is to endow personalized 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.

The rest of the paper is organized as follows. Section 2 introduces the notion of context and related work in this area. Our approach to contextual personalization is described in detail after that, including our underlying ontology-based personalization framework (Section 3.1), the proposed context representation model (Section 3.2), a mechanism to instantiate the model (Section 3.3), a method to filter user preferences by context (Section 3.4), and the final computation of a personalized retrieval function for preference-biased, context-sensitive result ranking (Section 3.5). Finally, some conclusions are given in Section 4.

## 2. 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 (Brown et al 1997). 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 the term is used (Edmon et al 1999). The term has a long history in diverse areas of computer science, namely in artificial intelligence, information retrieval, 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 remains still an open and challenging problem. Several researchers have tried 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). 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 (Abowd et al 1999).

This paper is concerned with exploiting contextual information and smoothly integrating it into the personalization of information retrieval. In this field, contextual information can be proven to be very helpful when dealing with information 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.

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 (Bharat 2000, Finkelstein et al 2002, Haveliwala 2003, Jones et al 2004, Lawrence et al 2000). 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 (Akrivas et al 2002). One common solution is the use of relevance feedback (Rocchio 1971). However, the effectiveness of relevance feedback is considered to be limited in real systems, basically because users are often reluctant to provide such information.

For this reason, implicit feedback has attracted greater attention recently (Campbell et al 1996, Kelly et al 2003, Ryen et al 2006). For a complex or difficult information need, 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, which documents the user has chosen to view, and even how a user has read a document. 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.

### 3. 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. 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, *contextualize* 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 the domain ontology. 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 ranges on a continuum scale instead, as will be seen, where the contextual weight of a preference decreases monotonically with the semantic distance between the preference and the context.

#### 3.1 Underlying Personalization Framework

The contextualization model presented here builds upon an ontology-based personalization framework developed in the aceMedia<sup>1</sup> project (Castells et al 2005). Building on ontology-based semantic structures and semantic metadata, the aceMedia 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 aceMedia retrieval system assumes that the items in a retrieval space  $\mathcal{D}$  are annotated with weighted semantic metadata which describe the meaning carried by the item, in terms of a domain ontology  $\mathcal{O}$ . That is, each item  $d \in \mathcal{D}$  is associated 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

---

<sup>1</sup> <http://www.acemedia.org>

the degree to which the concept  $x$  is important in the meaning of  $d$ .

The aceMedia 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 (numbers from 0 to 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 (a collection, a catalog section, a search result) are filtered and ranked in personalized ways. Further details of the aceMedia system can be found in (Castells et al 2005).

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, ontology standards, such as RDF and OWL, 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 items about 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 of Sicily can be assumed to like Palermo, through the transitive *locatedIn* relation.

### 3.2 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  $P : \mathcal{U} \rightarrow \mathcal{P}$  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  $C : \mathcal{U} \times \mathbb{N} \rightarrow \mathcal{C}$ , 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  $\Phi : \mathcal{P} \times \mathcal{C} \rightarrow \mathcal{P}$  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  $\mathcal{O}$ .

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 used in a personalized retrieval framework, so far we have not made any assumption on the type of application where the abstract model defined so far is to be implemented, so the 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  $\Phi$ , and the techniques to compute it. Once we define this, the activated user preferences in a given context will be given by  $\Phi(P(u), C(u,t))$ .

### 3.3 Building a Dynamic Retrieval Context

The model defined in the previous section 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 a drift of concepts over time, and a general approach towards achieving this follows. Right after each user's request, a request vector  $R(t) \in \mathcal{C}_{\mathcal{O}}$  is defined. This

vector may be defined as, for instance, 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 (Popov et al 2004). If the request is of the type “view document”,  $R(t)$  can be defined by the top-most 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 modalities.

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  $\xi$ , a real value in  $[0,1]$ . Consequently, at a given time  $t$ , we update  $C_x(t)$  as:

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

Although this may seem similar to a pseudo-relevance feedback strategy, here the context vector  $C(t)$  is not used to reformulate the query, but to focus the preference vector, as shown next.

### 3.4 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, the more in context a preference shall be considered. The semantic paths are explored by a form of Constraint Spreading Activation (Crestani 1997). 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 concepts represents the degree to which it belongs to each set (see figure 1).

After the context is 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  $CP_x = EP_x \cdot C_x$  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.  $CP(u,t)$ .

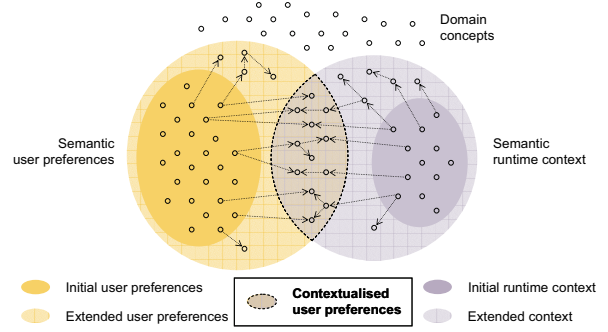


Figure 1. Contextual activation of semantic user preferences.

Note also that at this point we have achieved a contextual preference mapping  $\Phi$  as defined in Section 3.2, namely  $\Phi(P(u), C(u,t)) = CP(u,t)$ , where  $P(u) \models \Phi(P(u), C(u,t))$ , since  $CP_x(u,t) > P_x(u)$  only when  $EP_x(u)$  has been derived from  $P(u)$  through the constrained spreading expansion mechanism, and  $CP_x(u,t) < EP_x(u)$ .

### 3.5 Personalized Retrieval in Context

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

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

where  $M(d) \in [0,1]^{|\mathcal{O}|}$  is the semantic metadata concept-vector of the document, as explained in Section 3.1. 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  $f(x,y) = \lambda \cdot \bar{x} + (1 - \lambda) \cdot \bar{y}$ . The term  $\lambda$  is the personalization factor that shall determine the degree of personalization applied to the search result ranking. (Castells et al 2005) address the problem of how to set the value of  $\lambda$  dynamically.  $\bar{x}$  and  $\bar{y}$  denote the normalization of the score values  $x$  and  $y$ , which is needed before the combination to ensure e.g. that they range on the same scale (Fernández et al 2006). 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.

## 4. Conclusions

Context is an increasingly common notion in Information Retrieval. 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 (Bharat 2000, Finkelstein et al 2002, Haveliwala 2003, Jones et al 2004, Lawrence et al 2000).

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 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.

On another angle, 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.

## Acknowledgements

This research was supported by the European Commission (FP6-001765 – aceMedia), and the Spanish Ministry of Science and Education (TIN2005-06885). The content expressed is the view of the authors but not necessarily the view of the aceMedia project as a whole.

## References

- Abowd G. D., Dey, A. K., Brown, P. J., Davies, N., Smith, M., Steggles, P. 1999. Towards a Better Understanding of Context and Context-Awareness. In *Proc. of the 1<sup>st</sup> Int. Symp. on Handheld and Ubiquitous Computing*, LNCS Vol. 1707, 304-307. Berlin Heidelberg New York: Springer Verlag.
- Akrivas, G., Wallace, M., Andreou, G., Stamou, G., Kollias, S. 2002. Context-Sensitive Semantic Query Expansion. In *Proceedings of the IEEE International Conference on Artificial Intelligence Systems*, 109-104. Divnomorskoe, Russia: IEEE Computer Society Press.
- Bharat, K. 2000. SearchPad: Explicit capture of search context to support web search. In *Proc. of the 9<sup>th</sup> International World Wide Web Conf.* Amsterdam: ACM Press.
- Brown, P. J., Bovey, J., Chen, X 1997. Context-Aware Applications: From the Laboratory to the Marketplace. *IEEE Personal Communications* 4(5):58-64.
- Campbell, I., and Van Rijsbergen, C. J. 1996. The ostensive model of developing information needs. In *Proceedings of the 3<sup>rd</sup> International Conference on Conceptions of Library and Information Science*, 251-268.
- Castells, P., Fernández, M., Vallet, D., Mylonas, P., and Avrithis, Y. 2005. Self-Tuning Personalized Information Retrieval in an Ontology-Based Framework. In *Proc. of the 1<sup>st</sup> Int. Workshop on Web Semantics*, LNCS Vol. 3762, 977-986. Berlin Heidelberg New York: Springer Verlag.
- Crestani, F. 1997. Application of Spreading Activation Techniques in Information Retrieval. *Artificial Intelligence Review* 11:453-482.
- Edmonds, B 1999. The Pragmatic Roots of Context. In *Proc. of the 2<sup>nd</sup> International and Interdisciplinary Conference on Modeling and Using Context*, LNAI Vol. 1688, 119-132. Berlin Heidelberg New York: Springer Verlag.
- Fernández, M., Vallet, D., Castells, P. 2006. Probabilistic Score Normalization for Rank Aggregation. In *Proc. of the 28<sup>th</sup> European Conf. on Inf. Retrieval*, LNCS Vol. 3936, 553-556. Berlin Heidelberg New York: Springer Verlag.
- Finkelstein, L., Gabrilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolfman, G., Ruppin, E. 2002. Placing Search in Context: The Concept Revisited. *ACM Transaction on Information Systems* 20(1): 116-131.
- Gauch, S., Chaffee, J., and Pretschner, A. 2003. Ontology-based personalized search and browsing. *Web Intelligence and Agent Systems* 1(3-4):219-234.
- Haveliwala, T. H. 2003. Topic-Sensitive PageRank: A Context-Sensitive Ranking Algorithm for Web Search. *IEEE Trans. on Knowledge and Data Engineering* 15(4):784-796.
- Jones, G. J. F., and Brown, P. J. 2004. The Role of Context in Information Retrieval. In *Proc. of the SIGIR Information Retrieval in Context Workshop*. Sheffield, UK: ACM Press.
- Kelly D., Teevan, J. 2003. Implicit feedback for inferring user preference. *SIGIR Forum* 32(2):18-28.
- Lawrence S 2000. Context in Web Search. *IEEE Data Engineering Bulletin* 23 (3):25-32.
- Popov, B., Kiryakov, A., Ognyanoff, D., Manov, D., and Kirilov, A. 2004. KIM – A Semantic Platform for Information Extraction and Retrieval. *Journal of Natural Language Engineering* 10(3-4):375-392.
- Rocchio, J. 1971. Relevance feedback information retrieval. In Salton, G. (Eds.), *The SMART Retrieval System: Experiments in Automatic Document Processing*, 313-323. Englewood Cliffs, NJ: Prentice-Hall.
- White, R. W., Jose, J. M., Ruthven, I. 2006. An implicit feedback approach for interactive information retrieval. *Information Processing and Management* 42(1):166-190.