

Semantic Contextualisation in a News Recommender System

Iván Cantador^{1,2}, Pablo Castells²

¹ Department of Computing Science
University of Glasgow
Lilybank Gardens
G12 8QQ, Glasgow, UK
cantador@dcs.ac.gla.uk

² Departamento de Ingeniería Informática
Universidad Autónoma de Madrid
Campus de Cantoblanco
28049, Madrid, Spain

{ivan.cantador, pablo.castells}@uam.es

ABSTRACT

The elements that can be considered under the notion of context in a recommender system are manifold: user tasks/goals, recently browsed/rated items, computing platforms and network conditions, social environment, physical environment and location, time, external events, etc. Complementarily to these elements, we propose a particular notion of context for semantic content retrieval: that of semantic runtime context, which we define as the background topics under which activities of a user occur within a given unit of time. A runtime context is represented in our approach as a set of weighted concepts from domain ontologies, obtained by collecting the concepts that have been involved in user's actions (e.g., accessed items) during a session. Once the context is built, a contextual activation of user preferences is achieved by finding semantic paths linking preferences to context. In this paper, we present a user-centred study of our context-aware recommendation model using a news recommender system called News@hand. We analyse the strengths and weaknesses of our approach, and discuss the importance of contextualisation in a news recommendation scenario.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval – *information retrieval, retrieval models*. I.2.4 [Artificial Intelligence]: Knowledge Representation and Methods – *semantic networks*.

General Terms. Algorithms, Experimentation, Human Factors.

Keywords. Recommender systems, context modelling, ontology.

1. INTRODUCTION

With the advent of the Web, people nowadays not only have access to more worldwide news information than ever before, but can also obtain it in a more timely manner. Online newspapers present breaking news on their websites in real time, and users can receive automatic notifications about them via RSS feeds. Even with such facilities, further issues remain nonetheless to be addressed. The increasing volume, growth rate, ubiquity of access, and the unstructured nature of content challenge the limits of human processing capabilities.

It is in such scenario where recommender systems can do their most, by scanning the space of choices, and predicting the potential usefulness of news for each particular user, without explicitly specifying needs or querying for items whose existence is unknown beforehand. However, general common problems have not been fully solved yet. For example, typical approaches are domain dependent. Their models are generated from information gathered within a specific domain, and cannot be easily extended and/or

incorporated to other systems. Moreover, the need for further flexibility in the form of query-driven recommendations, and the consideration of contextual features during the recommendation processes are also unfulfilled requirements in most systems [1].

In this paper, we focus on the contextualisation of item recommendations. Specifically, we particularly define context as the background topics under which activities of a user occur within a given unit of time. Describing user preferences and item contents in terms of semantic concepts that belong to a number of domain ontologies, a runtime context is represented in our approach as a set of weighted concepts from such ontologies. This set is obtained by collecting the concepts that have been involved in the interaction of the user (e.g., accessed items) during a session. Once the context is built, a contextual activation of user preferences is achieved by finding semantic paths linking preferences to context. The perceived effect of contextualisation is that user interests that are out of focus, under a given context, are disregarded, and those that are in the semantic scope of the ongoing user activity are more considered for recommendation.

This context-aware recommendation model is integrated and evaluated in News@hand [6], a news recommender system. The results obtained from a preliminarily user-centred study show that semantic contextualisation improves the accuracy of personalised news recommendations, as well as increases the users' satisfaction on the news item suggestions.

The rest of the paper is organised as follows. Section 2 describes News@hand. Section 3 explains our semantic contextualisation approach, and Section 4 presents the conducted experiments. Finally, Section 5 gives some conclusions and future work lines.

2. NEWS@HAND SYSTEM

News@hand is a news recommender system that uses semantic technologies to provide several types of recommendations: driven by a concept-based query [8], personalised to a single user's profile [13], oriented to the interests shared by a group of users [7], combining content-based and collaborative recommendation techniques [5], and finally, considering the current topic context of the session.

Figure 1 shows a typical news recommendation page in News@hand. News items are classified into eight different sections: headlines, world, business, technology, science, health, sports, and entertainment. When the user is not logged in the system, he can browse any of the previous sections, but the items are listed without any personalised criterion. When the user is logged in, recommendation and user profile editing are enabled, and the user can browse the news according to his and others' preferences in different ways.

At the centre of the screen, for each news item, apart from its title, source, date, summary, image and link to the full article, additional information is shown. Those terms appearing in the item that are associated to semantic annotations of the contents, the user profile, and the current context are highlighted with different colours. Its global collaborative rating (a linear combination of the results obtained with a pure item-based collaborative filtering strategy, and a semantic multilayer hybrid recommendation technique [5]) is shown in a five-star scale, and two coloured bars indicate the relevance of the news item for the semantic user profile and the current user context, separately.

On the left side of the screen, the user can set the input parameters he wants for single or group-oriented recommendations: the consideration of preferences of the user, the user’s contacts, or all the users; the degree (weight) of relevance that the concepts of the semantic user profile and context should have in the recommendation algorithms; and multi-criteria conditions to be fulfilled by the user evaluations of the news articles to retrieve.



Figure 1. News@hand graphical user interface

2.1 Knowledge representation

Our recommendation approaches use a controlled and structured vocabulary to describe user preferences and news content. Working within an ontology-based personalisation framework [13], user preferences are represented as vectors $\mathbf{u}_m = (u_{m,1}, \dots, u_{m,K})$, where the weight $u_{m,k} \in [-1,1]$ measures the intensity of the interest of user u_m for concept $c_k \in \mathcal{O}$ (a class or an instance) in a domain ontology \mathcal{O} , K being the total number of concepts in the ontology. A positive weight indicates the user is interested in the concept, while a negative one reflects a user dislike. Similarly, items i_n in the retrieval space are assumed to be described (annotated) by vectors $\mathbf{i}_n = (i_{n,1}, \dots, i_{n,K})$ of concept weights $i_{n,k} \in [0,1]$ in the same vector-space as user preferences, measuring the relevance of the concepts in the items’ contents. Based on this common logical representation, measures of user interest for items can be computed by comparing preference and annotation vectors, and these measures can be used to prioritise, filter and rank contents.

The main benefits of a concept-based user profile representation versus common keyword-based approaches are the following:

- **Semantic richness.** Ontology concept-based preferences are more precise, and reduce the effect of the ambiguity caused by simple keyword terms. For instance, if a user states an interest for “java”, the system does not have further information to distinguish the programming language from the Pacific island.

A preference stated as `Island:Java` (this is read as the instance “Java” from the “Island” class) lets the system understand unambiguously the preference of the user.

- **Hierarchical representation.** Ontology concepts are represented in a hierarchical way, through different hierarchy properties, such as `subclassOf`, `instanceOf` or `partOf`. Parents, ancestors, children and descendants of a concept give valuable information about the semantics of the concept. For instance, the concept `leisure` might be highly enriched by the semantics of each leisure activity, which would be described by the taxonomy that the concept could subsume.
- **Inference.** Ontology standards support inference mechanisms that can be used to enhance recommendation, so that, for instance, a user interested in skiing, snowboarding and ice hockey can be inferred with a certain confidence to be globally interested in winter sports. Also, a user keen on USA can be assumed to like New York, through the `locatedIn` transitive relation, assuming that this relation had been seen as relevant for inferring previous user’s interests.

2.2 User preference expansion

To overcome sparsity in user profiles, we propose a semantic preference spreading mechanism which expands the initial set of preferences stored in user profiles through explicit semantic relations with other concepts in the ontology.

The activation of user preferences is based on an approximation to conditional probabilities. Let $u_{m,k} \in [-1,1]$ be the preference (interest/dislike) of the user u_m for the ontology concept $c_k \in \mathcal{O}$. The probability that $c_x \in \mathcal{O}$ is relevant for the user can be expressed in terms of the probability that c_x and each concept $c_y \in \mathcal{O}$ directly related to c_x in the ontology belong to the same topic, and the probability that c_y is relevant for the user. A similar formulation could be given for non-relevant concepts.

Let \mathcal{R} be the set of all relations in \mathcal{O} . The spreading strategy is based on weighting each semantic relation $r \in \mathcal{R}$ with a measure $w(r, c_x, c_y)$ that represents the probability that given the fact that $r(c_x, c_y)$ holds, c_x and c_y belong to the same topic. This is used for estimating the relevance of c_y when c_x is relevant for the user. With this measure, concepts are expanded through the relations of the ontology using a Constrained Spreading Activation (CSA) [9] mechanism over the semantic network defined by these relations. As a result, the initial set of concepts $\mathbf{P}_u = \{c_x \in \mathcal{O} | u_k \neq 0\}$ is extended to a larger vector \mathbf{EP}_u , which is computed as:

$$\mathbf{EP}_u[c_y] = \begin{cases} \mathbf{P}_u[c_y] & \text{if } \mathbf{P}_u[c_y] > 0 \\ R(\{\mathbf{EP}_u[c_x] \cdot power(c_x)\}_{c_x \in \mathcal{O}, r(c_x, c_y)}) & \text{otherwise} \end{cases}$$

where $power(c_x) \in [0,1]$ is a propagation power assigned to each concept c_x (1 by default), and

$$R(\mathbf{X}) = \sum_{S \subseteq \mathbb{N}_n} \left\{ (-1)^{|S|+1} \times \prod_{i \in S} x_i \right\},$$

having $\mathbf{X} = \{x_i\}_{i=0}^n, x_i \in [0,1]$. For further details about the previous formula, the reader is referred to [9].

2.3 Personalised recommendation

Assuming a semantic profile of user preferences has been obtained, either automatically or manually, our notion of personalised content retrieval is based on the definition of a matching algorithm that provides a personal relevance measure $pref(i_n, u_m)$ of an item i_n for a user u_m . This measure is set

according to the semantic preferences \mathbf{u}_m of the user and the semantic annotations \mathbf{i}_n of the item, and is based on the cosine function for vector similarity computation:

$$pref(i_n, u_m) = \cos(\mathbf{i}_n, \mathbf{u}_m) = \frac{\mathbf{i}_n \cdot \mathbf{u}_m}{\|\mathbf{i}_n\| \times \|\mathbf{u}_m\|}$$

The formula matches two weighted-concept vectors and produces a value in $[0,1]$. Values close to 0 are obtained when the two vectors are dissimilar, and indicate that user preferences negatively match the content metadata. On the other hand, values close to 1 indicate that user preferences significantly match the content metadata, which means a potential interest of the user for the item.

Personalisation should combine long-term preferences, based on past usage history, with contextualised predictions based on current user activities, as well as reactions to (implicit or explicit) user feedback to output. The incorporation of contextualised semantic preferences into the presented ontology-based personalised recommendation model is indeed the purpose of the work presented in the next sections.

3. SEMANTIC CONTEXT-AWARE RECOMMENDATION

Context is a difficult notion to capture in a recommender system, and the elements that can be considered under the notion of context are manifold: user tasks/goals, recently browsed/rated items, computing platforms and network conditions, social environment, physical environment and location, time, external events, etc. As representative examples, the reader is referred e.g. to [2], [4], [11], [12]. Complementarily to these, we propose a particular notion of context for semantic content retrieval: the semantic runtime context, which we define as the background topics \mathbf{C}_u^t under which activities of a user u occur within a given unit of time t . A runtime context is represented in our approach as a set of weighted concepts from a domain ontology \mathcal{O} . This set is obtained by collecting the concepts that have been involved in user's actions (e.g., accessed items) during a session. Similarly to [10], the context is built in such a way that the importance of concepts $c_k \in \mathcal{O}$ fades away with time (number of steps back when the concept occurred) by a decay factor $\xi \in [0,1]$:

$$\mathbf{C}_u^t[c_k] = \xi \cdot \mathbf{C}_u^{t-1}[c_k] + (1 - \xi) \cdot \mathbf{Req}_u^t[c_k]$$

where $\mathbf{Req}_u^t \in [0,1]^{|\mathcal{O}|}$ is a vector whose components measure the degree in which the concepts c_k are involved in the user's request at time t . This vector can be defined in multiple ways, depending on the application: a query concept-vector (if a request is expressed in term of a concept-based search query), a concept vector containing the most relevant concepts in a document (if a request is a "view document" request), the average concept-vector corresponding to a set of items marked as relevant by the user (if a request is a relevance feedback step), etc. The decay factor ξ establishes the number of action units in which a concept is considered as in the current semantic context, i.e., how fast a concept is "forgotten" by the system when recommendations have to be made.

Once the context is built, a contextual activation of preferences is achieved by finding semantic paths linking preferences to context. These paths are made of existing relations between concepts in the ontology, following the CSA technique explained in Section 2.2. This process can be understood as finding an intersection between user preferences and the semantic context, where the final computed weight of each concept represents the degree to which it belongs to each set (Figure 2). The perceived effect of contextualisation is that user interests that are out of focus, under a given

context, are disregarded, and those that are in the semantic scope of the ongoing user activity are considered for recommendation.




Figure 2. Expansion and contextualisation of user preferences

After the semantic user profile \mathbf{P}_u^t and context \mathbf{C}_u^t are propagated through the ontology relations, a combination of their expanded versions \mathbf{EP}_u^t and \mathbf{EC}_u^t is exploited for making context-aware personalised recommendations using the following expression:

$$\begin{aligned} pref_c(i, u) &= \lambda \cdot pref(i, \mathbf{EP}_u^t) + (1 - \lambda) \cdot pref(i, \mathbf{EC}_u^t) \\ &= \lambda \cdot \cos(\mathbf{i}, \mathbf{EP}_u^t) + (1 - \lambda) \cdot \cos(\mathbf{i}, \mathbf{EC}_u^t) \end{aligned}$$

where $\lambda \in [0,1]$ measures the strength of the personalisation component with respect to the current context. This parameter could be manually established by the user, or dynamically adapted by the system according to multiple factors, such as the current size of the context, the automatic detection of a change in the user's search focus, etc.

4. EXPERIMENTS

4.1 Knowledge base

In News@hand, ontologies are populated with semantic concepts associated to noun terms extracted from the news contents. These terms are categorised as common nouns (e.g., actor) and proper nouns (e.g., Brad Pitt). Common nouns are easily processable because their corresponding concepts can be found in English dictionaries like WordNet. Proper nouns may require a more complex processing. In order to infer their concepts, general multi-domain knowledge is needed. We propose to extract that information from Wikipedia, a multilingual, open-access, free encyclopaedia on the Internet. Its articles describe a number of different types of entities: people, places, companies, etc., providing descriptions, references, and even images about the described entities. In addition to the above elements, every Wikipedia article contains a set of categories that give an idea of the meaning of the associated concept. We have implemented an automatic mechanism that creates ontology instances using, among other things, the Wikipedia categories of the terms [6]. The basic idea of our proposal is to match the categories of an entity with classes of the ontologies, and then link the entity with the matched ontology class that is most "similar" to the entity categories.

A total of 17 ontologies have been used for the current version of the system. They are adaptations of the IPTC ontology [6], which contains concepts of multiple domains such as education, politics, religion, science, technology, business, health, entertainment, sports, weather, etc. A total of 137,254 Wikipedia entries were used to populate 744 ontology classes with 121,135 instances.

4.2 Item annotations

News@hand periodically retrieves news items from the websites of well-known media sources, such as BBC, CNN, The New York Times, and The Washington Post. These items are obtained via RSS, and contain information of published news articles: their title, summary of contents, publication date, hyperlinks to full texts and related images. The system analyses and annotates the textual information (title and summary) of the RSS feeds with concepts that exist in the domain ontologies and have been previously indexed.

Using a set of NLP tools [3], an annotation module removes stop words, and extracts relevant (simple and compound) terms, categorised according to their Part of Speech: nouns, verbs, adjectives, etc. Then, nouns are morphologically compared with the names of the classes and instances of the domain ontologies. The comparisons are done using an ontology index, and according to fuzzy metrics based on the Levenshtein distance. For each term, if similarities above a certain threshold are found, the most similar concepts are chosen and added as annotations of the news items. After all the annotations are created, a TF-IDF technique computes and assigns weights to them. For more details, see [6].

We have run our semantic annotation approach on a set of 9,698 news items daily retrieved during two months. The ontological Knowledge Base (KB) from which we obtained the semantic concepts appearing in the annotations is the one explained in Section 4.1. A total of 66,378 annotations were created.

4.3 Experimental setting

We conducted an experiment to evaluate the precision of the personalised and context-aware recommendation functionalities available in News@hand, and to investigate the influence of each mechanism in the integrated system, measuring the precision of the recommendations when a combination of both models is used.

The experiment was done with 16 subjects, recruited in our department. They were graduate students and lecturers. The experiment consisted of two phases, each composed of two tasks.

- In the first phase, only the personalisation module was active, and its tasks were different in having the semantic expansion enabled or disabled.
- In the second phase, the contextualisation and semantic expansion functionalities were active. On its second task we also enabled the personalised recommendations.

4.3.1 Search tasks

A task was defined as finding out and evaluating those news items that were relevant to a given goal. Each goal was framed in a specific domain. We considered three domains: telecommunications, banking and social care issues. For each domain, a user profile and two search goals were manually defined (see below). Table 1 shows a summary of the involved tasks.

Table 1. Summary of the performed search tasks

Domain	Section	Query	Task goal
Telecom	World	Q _{1,1} pakistan	Media: TV, radio, internet
	Entertainment	Q _{1,2} music	Software piracy, illegal downloads, file sharing
Banking	Business	Q _{2,1} dollar	Oil prices
	Headlines	Q _{2,2} fraud	Money losses
Social care	Science	Q _{3,1} food	Cloning
	Headlines	Q _{3,2} internet	Children, young people, child safety, child abuse

To simplify the searching tasks, limiting the number of items to be evaluated by the subjects, they were defined for pre-established sections and queries. For example, the task goal of finding news items about software piracy, illegal downloads and file sharing, Q_{1,2}, was reduced to evaluate those articles existing in Entertainment section that were retrieved from the query “music”. The configuration and assignment of the tasks were uniformly set according to the following principles:

- A user did not repeat a query during the experiment.
- The domains were equally covered by each experiment phase.
- A user had to manually define a user profile once in the experiment.

4.3.2 User profiles

The user profile editor of News@hand allows the users to manually create and update their semantic preferences. An ontology browser lets explore the ontology hierarchies, easily search for concepts through on-line auto-complete widgets, and add selected concepts into the profile assigning weights to them

As mentioned before, fixed user profiles were used for each domain. Some of them were common predefined profiles, and others were created by the users during the experiment using the profile editor. In addition, some tasks were done with user profiles containing concepts belonging to all the three domains. Each domain was described with 6 semantic concepts, appearing in a significant number of item annotations. Note that each domain may be described by concepts belonging to different ontologies, and may be covered with news items of different news sections.

Analogously to the predefined user profiles, those manually created by the evaluators contained concepts of the above three domains. However, in this case, the evaluators were free to select their preferences from concepts available in the entire system KB. No restriction was placed on the number, type (classes or instances) and ontology of the concepts. For instance, in *Telecommunication* domain, 55 preferences were declared using 30 different semantic concepts, producing an average of 3.4 preferences per user. On average, each profile contained 3.2 preferences of each domain.

4.3.3 Evaluation protocol

The objective of the two tasks performed in the first experiment phase was to assess the importance of activating the semantic expansion in our recommendation models. The following are the steps the users had to do in these tasks.

1. Launch the query with the personalisation deactivated.
2. Rate the top 15 news items. The allowed rating values were: 1 if the item was not relevant to the task goal, 2 if the item was relevant to the task goal, and 3 if the item was relevant to the task goal and the user profile. These ratings are considered as our *baseline* case.
3. Launch the query with the personalisation activated (and the semantic expansion enabled/disabled depending on the case).
4. Rate the new top 15 items as explained in 2. If an item had previously been rated, rate it again with the same value.

The objective of the two tasks performed in the second experiment phase was to assess the quality of the results when the contextualisation functionality is activated and combined with personalisation. The steps done in this phase were:

1. Launch the query with the contextualisation deactivated.
2. Rate the top 15 news items as explained before, and evaluate as relevant (clicking the title) the first two items which were related to the task goal. Doing this the current semantic context is updated.
3. Launch the query with the **contextualisation activated** (semantic expansion enabled, and personalisation enabled/disabled depending on the case).
4. Rate again the top 15 items as explained in 2. If an item had previously been rated, rate it again with the same value.

4.4 Results

4.4.1 Analytical evaluation

Once the two evaluation phases were finished, we computed the precision values for the top $N = 5, 10, 15$ news items as follows:

$$P@N = \frac{\#\{\text{relevant items in the top } N \text{ news items}\}}{N}$$

Figure 3 shows the average results for the 16 users, taking into account those items evaluated as relevant to the task goal, and also the user profile.

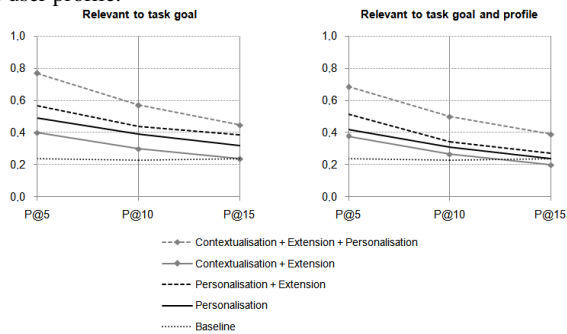


Figure 3. Average precision values for the top 5, 10 and 15 news items, taking into account those items evaluated as relevant to the task goal and the user profile

In both cases, the recommendation models outperformed the baseline, especially for the 5 top items. The $P@5$ values increased from 20% of the baseline case to almost 40% and 50% when contextualisation and personalisation functionalities were enabled. The semantic expansion seemed to be an essential component within the recommendation processes. It provided an improvement of 10% in the personalisation precision. Finally, the combination of personalised and context-aware recommendations (plus semantic expansion) gave the best results, achieving a $P@5$ value of 80%.

4.4.2 User questionnaires

Apart from the computation of the precision values, we also asked the evaluators to fill structured questionnaires about their news consumption preferences and habits, and to provide comments and suggestions about the system. The most remarkable feedback we obtained can be summarised in the following points:

- The contextualisation of recommendations is a useful functionality. The users noticed and positively assessed how items relevant to the current search goal move up to the top positions when the context-aware recommender was activated.
- A disambiguation mechanism should be included within the annotation process. The users found out semantic annotations whose terms appeared in their profiles but having different meanings. This not only worsened the generated recommendations, but also the users' evaluations.
- A collaborative approach to enrich the semantic profiles may be beneficial. Several users declared some preferences assuming that related ones (e.g., synonyms) were going to be implicitly taken into account. A mechanism to exploit co-occurrences among preferences of different users could be useful to automatically add related concepts into the profiles.
- The incorporation of a user preference recommender would be helpful. Despite the facilities offered by the ontology browser and the auto-complete concept search boxes of News@hand, several users missed the fact of having concept suggestions (e.g., in the form of "related preferences are...") when they had to create their profiles.

5. CONCLUSIONS AND FUTURE WORK

We have presented and evaluated semantic personalisation and contextualisation models in a news recommender system. The personalised recommendations helped users to find relevant news articles, and a semantic expansion of user preferences eased the matching between user and item profiles, improving precision values for the top suggested items, and mitigating the cold-start and sparsity problems. The incorporation of contextualisation within the personalisation mechanism speeded up the discovery of items related to current search goals, and was highly appreciated by users.

The experiments also provided us the opportunity of getting feedback from users about the system functionalities and outputs. Among other issues, they showed the need of incorporating a disambiguation step in the semantic annotation process, and addressing of the non-diversity problem, as very similar news items were presented closely. Moreover, they suggested additional improvements in the profile editor, such as the integration of a real-time preference recommender taking into account concepts similar to the ones already introduced (synonyms, co-occurrences, etc.).

6. ACKNOWLEDGEMENTS

This work was supported by the Spanish Ministry of Science and Innovation (TIN2008-06566-C04-02) and the Ministry of Industry, Tourism and Commerce (CENIT-2007-1012).

7. REFERENCES

- [1] Adomavicius, G., Sankaranarayanan, R., Sen, S., Tuzhilin, A. 2005. *Incorporating Contextual Information in Recommender Systems using a Multidimensional Approach*. ACM TIS, 23(1), 103–145.
- [2] Ahn, J., Brusilovsky, P., Grady, J., He, D., Syn, S. Y. 2007. *Open User Profiles for Adaptive News Systems: Help or Harm?* 16th Intl. World Wide Web Conference (WWW 2007), 11–20.
- [3] Alfonseca, E., Moreno-Sandoval, A., Guirao, J. M., Ruiz-Casado, M. 2006. *The Wraetlic NLP Suite*. 5th International Conference on Language Resources and Evaluation (LREC 2006).
- [4] Billsus, D., Pazzani, M. J. 2000. *User Modeling for Adaptive News Access*. User Modeling and User-Adapted Interaction, 10(2), 147–180.
- [5] Cantador, I., Bellogín, A., Castells, P. 2008. *A Multilayer Ontology-based Hybrid Recommendation Model*. AI Communications, 21(2–3), 203–210.
- [6] Cantador, I., Bellogín, A., Castells, P. 2008. *News@hand: A Semantic Web Approach to Recommending News*. 5th Intl. Conf. on Adaptive Hypermedia and Adaptive Web-Based Systems (AH 2008), 279–283.
- [7] Cantador, I., Castells, P. 2009. *Extracting Multilayered Semantic Communities of Interest from Ontology-based User Profiles: Application to Group Modelling and Hybrid Recommendations*. Computers in Human Behavior. To appear.
- [8] Castells, P., Fernández, M., & Vallet, D. 2007. *An Adaptation of the Vector-Space Model for Ontology-based Information Retrieval*. IEEE TKDE, 19 (2), 261–272.
- [9] Crestani, F. 1997. *Application of Spreading Activation Techniques in Information Retrieval*. AI Review, 11(6), 453–482.
- [10] Middleton, S. E., Shadbolt, N. R., Roure, D. D. 2004. *Ontological User Profiling in Recommender Systems*. ACM Transactions on Information Systems, 22(1), 54–88.
- [11] Räck, C., Arbanowski, S., Steglich, S. 2006. *Context-aware, Ontology-based Recommendations*. International Symposium on Applications and the Internet Workshops (SAINTW 2006), 98–104.
- [12] Sujiyama, K., Hatano, K., Yoshikawa, M. 2004. *Adaptive Web Search Based On User Profile Constructed Without Any Effort From Users*. 13th Intl. World Wide Web Conference (WWW 2004), 675–684.
- [13] Vallet, D., Castells, P., Fernández, M., Mylonas, P., & Avrithis, Y. 2007. *Personalized Content Retrieval in Context Using Ontological Knowledge*. IEEE TCSVT 17 (3), 336–346.