

Ontology-based Personalised and Context-aware Recommendations of News Items

Iván Cantador, Alejandro Bellogín, Pablo Castells

Escuela Politécnica Superior, Universidad Autónoma de Madrid, Madrid, Spain
{ivan.cantador, alejandro.bellogin, pablo.castells}@uam.es

Abstract

News@hand is a news recommender system which makes use of semantic technologies to provide several on-line news recommendation services. News contents and user preferences are described in terms of concepts appearing in a set of domain ontologies. Based on the similarities between item descriptions and user profiles, and the semantic relations between concepts, content-based and collaborative recommendation models are supported by the system. In this paper, we focus on those models which are centred on a single user. Specifically, we evaluate a model that personalises the order in which news articles are shown to the user according to his long-term interest profile, and other model that reorders the news items lists taking into account the current semantic context of interest of the user. The combination of those models is investigated showing significant improvements on the experimental tasks performed.

1. Introduction

Thanks to the Web, people not only have access to more worldwide news information than ever before, but also they can obtain it in a more timely manner. On-line newspapers present breaking news on their websites in real time, and users can receive automatic notifications of them via RSS feeds. However, the increasing volume, growth rate and ubiquity of access to the contents challenge the limits of human processing capabilities. It is in such scenario where Recommender Systems (RS) can do their most, by scanning the space of choices, and predicting the potential usefulness of news items for each particular user, without explicitly specifying needs or querying for items whose existence is unknown beforehand.

Content-based and Collaborative Filtering (CF) techniques suffer from own limitations, and hybrid recommendation models have been proposed to address them by means of joined mechanisms [10][14]. However, general common problems have not been fully solved yet, and further investigation is needed. For example, typical approaches are domain dependant. 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 or group-

oriented recommendations, and the consideration of contextual features during the recommendation processes are also unfulfilled requirements in most systems.

In this work, we present News@hand, a system that makes use of semantic technologies to recommend news. The system supports different recommendation models for single and multiple users. The exploitation of meta-information in the form of ontologies that describe the recommended items and user profiles in a general, portable way, along with the capability of inferring knowledge from the semantic relations defined in the ontologies, are the key aspects of the proposals.

The recommendation models of News@hand have already been evaluated in different works [5][6][7]. However, previous experiments were conducted for each model in an isolated way. Here, we present our first attempt to jointly evaluate some of those models, integrated and combined within the system. Specifically, we examine two models centred on a single user: one that personalise the ranking of news contents according to the user's profile, and other that takes into consideration the current semantic context of interest of the user.

The rest of the paper is organised as follows. Section 2 describes existing approaches to recommending news. Section 3 presents the architecture and web-based user interface of News@hand. Its recommendation models are introduced in section 4. Experiments with its personalised and context-aware recommendation functionalities are shown in section 5. Finally, section 6 outlines some conclusions and future research lines.

2. News recommender systems

2.1. Content-based recommender systems

In content-based approaches, articles are suggested according to a comparison between their contents and the user profiles, the latter containing information about the users' content-based tastes.

NewsDude [3] is a personal news agent that uses a separate model for short-term and long-term interests. To determine short-term recommendations, news stories are described in terms of TF-IDF vectors, and are provided to a learning module based on NN algorithm. To establish long-term recommendations, news stories are represented as Boolean feature vectors, indicating the presence or absence of specific words, and are presented to a Bayesian learning module. News4U [12] is a system where

articles from a variety of online news sources are used to create a personalised newspaper. The user can decide which news sources to include, and can choose from a list of topics those he is interested in. YourNews [2] is a personalised news system which allows users to view and edit their profiles. It periodically gathers new articles from RSS feeds, passing them to an indexing module to build an index based on title, description, and content, and create and store TF-IDF term vectors of the articles. Users are provided a number of different news rankings according to specific selections of a topic, a time period (short and long-term preferences), and a type of view (recent and recommended news).

The content-based features can be combined with additional information, such as implicit behaviour of the user or explicit relevance feedback. PENS [16] is a personalised news system designed to provide adaptation to user location, user navigation history, and different user devices. NewT (News Tailor) [15] is a system, where the user can provide positive or negative feedback on articles, parts of an article, authors or sources, and this feedback is used to update the corresponding agent. Daily Learner [4] is an adaptive news service in which a user first chooses categories he wants to receive news about. Based on this profile, the system delivers those stories that best match the user's interests. Then, the user explicitly provides feedback using four rates (interesting, not interesting, more information, already known). Short-term interests are determined by analysing the N most recently rated stories.

2.2. Collaborative filtering systems

In CF systems, news items are suggested to a particular user according to the articles previously evaluated by other users. In general, users evaluate the articles submitting ratings. The items that were rated highly by the user's nearest neighbours and were not rated by the user are finally recommended.

GroupLens project [13] is one of the most referenced CF approaches. Its Netnews recommender is based on a client/server architecture, where users and Netnews are clustered according to the existing news groups, and implicit ratings are built measuring the time the users spent reading the articles. Personalised Google News [11] generates recommendations with three techniques: CF using MinHash clustering, Probabilistic Latent Semantic Indexing, and co-visitation counts. These techniques are combined using a linear model providing a scalable recommendation framework. The system also presents suggestions to users based on their click history and the click history of the community.

2.3. Hybrid recommender systems

Hybrid techniques combine content-based and CF strategies, mitigating inherent limitations of either paradigm.

NewsWeeder [14] is a Netnews filtering system in which the user can have access to news through virtual newsgroups, where lists of articles have been selected and ranked. The user must rate each article to have access to the following one with a rate from one to five. The system uses the collected rating information to learn a new model of the user's interests. Tango [10] presents an online newspaper recommender which bases a prediction on a weighted average of content-based and collaborative predictions. Articles are described as a set of keywords and the newspaper sections they belong to. Profiles are divided into segments corresponding to the newspaper sections. Each segment contains a set of explicit ratings and keywords given by the user, and a list of implicit keywords obtained from the highly rated articles.

3. News@hand

3.1. Architecture

As a hybrid recommender system, News@hand (<http://newsathand.net>) combines content features and collaborative information to make news suggestions. However, as opposite to previous systems, it uses a controlled and structured vocabulary to describe the news texts and user preferences. News items and user profiles are represented in terms of concepts appearing in domain ontologies, and semantic relations among those concepts are exploited to enrich the above representations, and are incorporated within the recommendation processes.

An ontology-based knowledge representation is less ambiguous than a keyword-based or item-based model, providing an adequate grounding for the representation of coarse to fine-grained features. Furthermore, ontology standards, such as RDF and OWL, support inference mechanisms that can be used to enhance content retrieval [9]. A user interested in *natural disasters* (superclass of *hurricane*) is also recommended items about *hurricanes*. Inversely, a user interested in *skiing* and *snowboarding* can be inferred with a certain confidence to be interested in other *winter sports*. Similarly, a user fascinated about the life of *actors* can be recommended items in which the name of *Brad Pitt* appears, due to that person could be an instance of the class *actor*. Also, a user keen on *Spain* can be assumed to like *Madrid*, through the *locatedIn* transitive relation between these concepts. Another advantages of this knowledge representation are its portability, thanks to the XML-based standards, the domain independency of the subsequent recommendation algorithms, and the multi-source nature of the proposal (different types of media could be semantically annotated: texts, images, videos).

Figure 1 shows how ontology-based item descriptions and user profiles are created in News@hand. Like other systems [2][12][16], news are automatic and periodically retrieved from several on-line news services via RSS feeds. The title and summary of the retrieved news are

then annotated with concepts of the domain ontologies. Thus, for example, all the news about actors, actresses, and similar terms might be annotated with the concept “actor”. Similarly to [2][3], a TF-IDF technique is applied to assign weights to the annotated concepts.

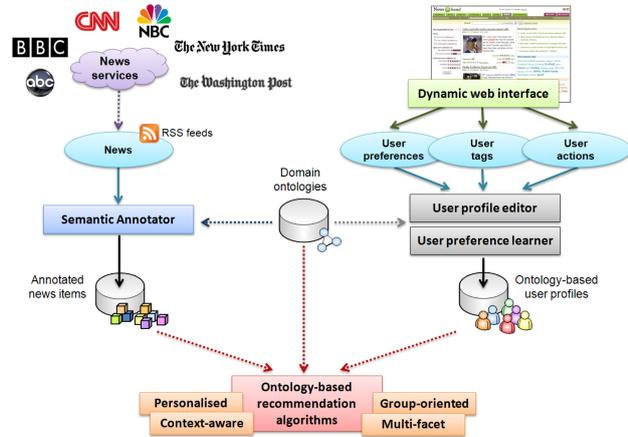


Figure 1. Architecture of News@hand

Using AJAX technology, a dynamic web interface allows the system to automatically store all the users’ inputs, analyse their behaviour with the system, update their preferences, and adjust the news recommendations in real time. As done in [10], explicit and implicit user preferences are taken into account, via manual preferences, tags and ratings, and via automatic learning from the users’ actions.

Finally, deriving benefit from the semantically annotated news items, the defined ontology-based user profiles, and the knowledge represented by the domain ontologies, a set of recommendation algorithms are executed. News@hand offers personalised, context-aware, group-oriented and multi-facet recommendations.

3.2. Web-based user interface

Figure 2 shows a screenshot of a news page in News@hand. The news items are classified in 8 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 news items are listed without any personalisation criterion. He can only sort them by their publication date, source or level of popularity (i.e., according to a classic rating-based CF mechanism [13]). On the other hand, when the user is logged in, recommendation and profile edition functionalities are enabled, and the user can browse the news according to his and others’ preferences in different ways. Like in other approaches [2][3][4], short and long term preferences are considered. As done in [11], click history is used to define the short term concepts, and similarly to [16], the resultant rankings can be adapted to the current context of interest.

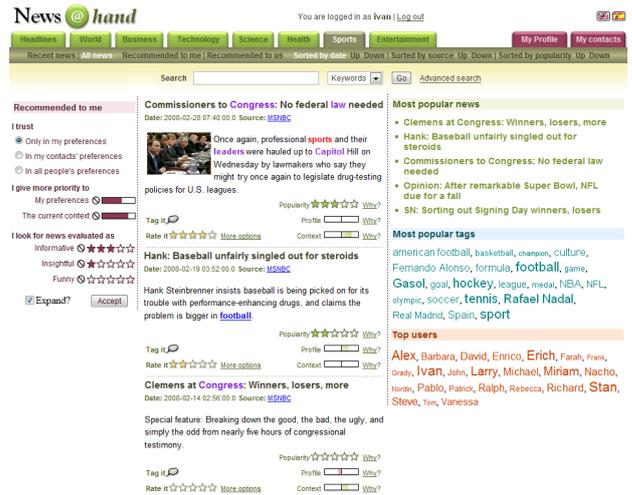


Figure 2. News recommendations page of News@hand

In the middle of the screen, apart from its title, source, date, summary, image, and link to the full article, additional information is shown for each item. Those terms associated to content annotations, user preferences, and context concepts are highlighted with different colours. The CF rating is shown in a 5-star scale, and two coloured bars indicate the relevance of the item for the user profile and the context. The user can view, and add comments, tags and ratings to the article, similarly to [14][15]. On the left side of the screen, he can set parameters for single or group recommendations: the consideration of preferences of his profile, his contacts, or all the users, the relevance degree that concepts of the profile and context should have, and multi-criteria conditions to be fulfilled by others’ evaluations.

3.3. Item annotation

News@hand (Figure 3) periodically retrieves news items from the websites of well-known news media. These items are obtained via RSS feeds, and contain information of published news articles: their title, summary of the contents, publication date, hyperlinks to the full texts and related on-line images.

The system analyses and automatically annotates the textual information (title and summary) of the RSS feeds. Using a set of NLP tools [1], an annotation module removes stop words and extracts relevant (simple and compound) terms, categorised according to their part of speech: nouns, verbs, adjectives or adverbs. Then, nouns are morphologically compared with the names of the classes and instances of the domain ontologies. The comparisons are done using an index created with Lucene (<http://lucene.apache.org/>), and according to fuzzy metrics based on the Levenshtein distance. For each term, if similarities above a certain threshold are found, the most similar semantic concept is chosen and added as an annotation of the item. After all annotations are created, a TF-IDF technique computes and assigns weights to them.

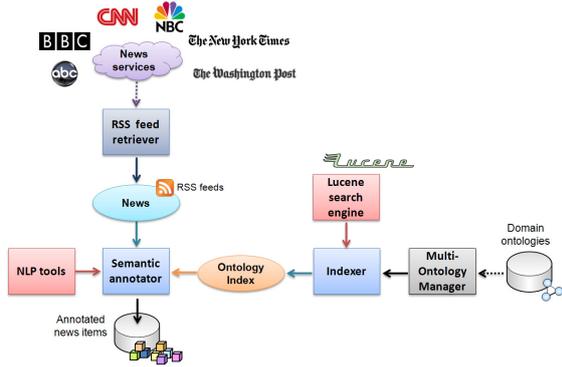


Figure 3. Item annotation mechanism of News@hand

3.4. Ontology population

As mentioned before, News@hand automatically annotates news items with semantic concepts that appear in the textual contents, and exist in the system domain ontologies. These concepts can be ontology classes (such as *actor*) or ontology instances (such as *Brad Pitt*, which might belong to the class *actor*). The ontology class hierarchy (i.e., the conceptual category schema) has to be fixed before the execution of the annotation process, and does not change when that process has finished. If modifications over the class hierarchy are done, the annotation mechanism has to be executed again. The ontology instances, however, can be created and associated to classes afterwards. In fact, the ontology population (i.e., the transformation of unstructured, semi-structured and structured data into ontology instances) is performed when new contents are retrieved from the Web.

Ontology population is an open research challenge to the Semantic Web community. In News@hand, a novel and simple mechanism to automatically populate ontologies using semantic information extracted from Wikipedia has been implemented. This process is explained in detail in [8]. Here, we only list its main steps. For each new term appearing in a news item, we:

- Search for the term in Wikipedia and retrieve the corresponding wiki page.
- Extract the Wikipedia categories of the term from the retrieved page. For example, *action film actors*, *american film actors*, etc. for *Brad Pitt*.
- Apply a morphologic matching algorithm between the term name and categories, and the names of the ontology classes.
- Select the ontology class whose name is most similar to the term name and/or categories.
- Create an instance for the term in the selected class.

4. Personalised recommendation models

4.1. Semantic content-based recommendations

Our notion of content retrieval is based on 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 of the user and the semantic annotations of the item, and is based on a cosine-based vector similarity $\cos(i_n, u_m)$. The measure can be combined with query-based scores $\text{sim}(i_n, q)$ and semantic context, to produce combined rankings [9][17].

In this scenario, and similarly to other RS, a main problem arises: the existence of sparsity in the user profiles. In general, users are not willing to spend time describing their preferences to the system, even less to assign weights to them. On the other hand, applications where an automatic preference learning algorithm is applied tend to recognise the main characteristics of user preferences. To overcome this problem, we propose a preference spreading mechanism, which expands the initial set of preferences stored in user profiles through explicit semantic relations with other concepts in the ontology (Figure 4). Our approach is based on Constrained Spreading Activation (CSA), and is self-controlled by applying a decay factor to the intensity of preference each time a relation is traversed [5][17].

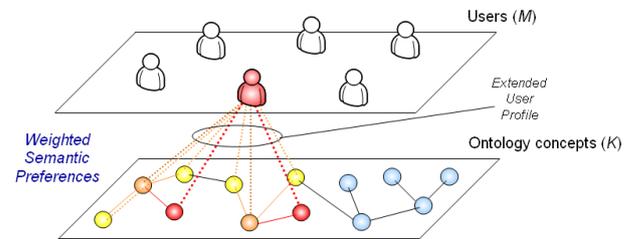


Figure 4. Semantic expansion of user preferences

4.2. Semantic context-aware recommendations

Context is a difficult notion to capture in a software system, and the elements that have been considered in the literature under the concept of context are manifold: user tasks/goals, computing platform, network conditions, social and physical environment, location, time, text around a word, visual content of a graphic region, etc.

Complementarily to the ones mentioned, we propose a particular notion useful in semantic content retrieval: that of semantic runtime context, which we define as the background topics under which user activities occur within a given unit of time. A runtime context is represented in our approach [17] as a set of weighted concepts from the domain 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. The context is built in such a way that the importance of concepts fades away with time (number of accesses back when the concept was referenced) by a decay factor.

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 ontologies, following the CSA technique introduced in section 4.1. 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 5). The perceived effect of contextualisation is that user interests that are out of focus, under a given context, are disregarded.

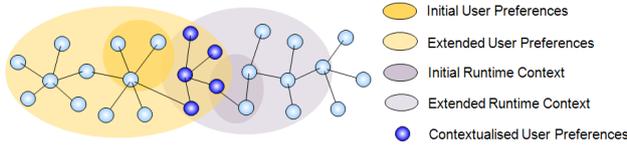


Figure 5. Semantic preference contextualisation

5. Experiments

We present an experiment conducted to evaluate the precision of the personalisation and context-aware recommendation functionalities of News@hand. With this experiment we also wanted 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.

5.1. News item database

For two months we were daily gathering RSS feeds. A total of 9,698 news items were stored. With this dataset, we run our semantic annotation mechanism introduced in section 3.3, and a total of 66,378 annotations were obtained. Table 1 shows a summary of the average number of annotations per news item generated with our system. We asked 20 graduate students of our department to evaluate 80 different news items from each of the 8 topic sections of News@hand, specifying whether the annotations of each item were correct or not. The annotation accuracies for each topic are presented in the table. An average accuracy of 74.8% was obtained.

Section	#news items	#annotations	#annotations /item	Avg. accuracy
Headlines	2,600	18,210	7	71.4
World	2,200	17,767	8	72.7
Business	1,739	13,090	8	79.2
Technology	303	2,154	7	76.3
Science	346	2,478	7	74.1
Health	803	4,874	6	73.1
Sports	603	2,453	4	75.8
Entertainment	1,044	5,343	5	76.0
	9,638	66,369	7	74.8

Table 1. Average number of annotations per news item, and average annotation accuracy values

5.2. Knowledge base

A set of 17 ontologies are used by the current version of the system. They are adaptations of the IPTC ontology,

(http://nets.ii.uam.es/iptc-kb_v01.zip) which contains concepts of multiple domains such as education, culture, politics, religion, science, technology, business, health, entertainment, sports, weather, etc. They have been populated with concepts appearing in the gathered news items using semantic information from Wikipedia, and applying the population mechanism mentioned in section 3.4. A total of 137,254 Wikipedia entries were used to populate 744 ontology classes with 121,135 instances.

Table 2 summarises the characteristics of the knowledge base used by our annotation mechanism, and shows the population accuracy values obtained for each ontology (see [8] for more details). The average accuracy of class assignments was 69.9%, and the average accuracy of ontology assignments arose 84.4%.

Ontology	#classes	#instances	Avg. accuracy
<i>arts, culture, entertainment</i>	87	33,278	78.7 / 93.3
<i>crime, law, justice</i>	22	971	62.7 / 73.3
<i>disasters, accidents</i>	16	287	74.7 / 84.0
<i>economy, business, finance</i>	161	25,345	69.3 / 80.0
<i>education</i>	20	3,542	75.5 / 76.7
<i>environmental issues</i>	41	20,581	72.0 / 85.3
<i>health</i>	26	1,078	65.3 / 89.3
<i>human interests</i>	6	576	64.0 / 84.0
<i>labour</i>	6	133	70.7 / 78.7
<i>lifestyle, leisure</i>	29	4,895	72.0 / 90.7
<i>politics</i>	54	3,206	60.0 / 81.3
<i>religion, belief</i>	31	3,248	84.0 / 90.7
<i>science, technology</i>	50	7,869	68.0 / 86.7
<i>social issues</i>	39	8,673	70.7 / 85.3
<i>sports</i>	124	5,567	72.0 / 86.7
<i>unrests, conflicts, wars</i>	23	1,820	61.3 / 80.0
<i>weather</i>	9	66	69.7 / 89.5
	744	121,135	69.9 / 84.4

Table 2. Number of ontology classes and instances in the KB, and average accuracy values of the population process

5.3. Experimental setup

Sixteen members of our department were requested to participate. In this case, they were undergraduate/graduate students, and lecturers. The experiment consisted of two phases, and each phase was composed of two different 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. More details are given in the next subsection.

Search tasks

In the experiment, a task was defined as finding 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 defined. Table 3 shows a summary of the involved tasks.

Profile	Section	Query	Task goal
1 Telecom	World	Q _{1,1} pakistan	News about media: TV, radio, Internet
	Entertainment	Q _{1,2} music	News about software piracy, illegal downloads, file sharing
2 Banking	Business	Q _{2,1} dollar	News about oil prices
	Headlines	Q _{2,2} fraud	News about money losses
3 Social care	Science	Q _{3,1} food	News about cloning
	Headlines	Q _{3,2} internet	News about children, young people, child safety, child abuse

Table 3. Summary of the search tasks performed in the experiment

To simplify the searching tasks (to limit the number of news items to be read by the users), they were defined for a pre-established section and query. Hence, 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 with the query “music”.

Table 4 shows the tasks performed by the 16 users. The configuration and assignment of the tasks was set according to the following principles:

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

For each phase, the combination of personalised and context-aware recommendations was established as a linear combination of their results using two weighs w_p , w_c in $[0,1]$:

$$score(i_n, u_m) = w_p \cdot pref(i_n, u_m) + w_c \cdot pref(i_n, u_m, context)$$

In the personalisation phase, the contextualisation was disabled (i.e. $w_c=0$). Its first tasks were performed without semantic expansion, and its second tasks had the semantic expansion activated. In the contextualisation phase, w_c was set to 1 and the expansion was enabled. Its first tasks were done without personalisation ($w_p=0$), and its second tasks were a bit influenced by the user profiles ($w_p=0.5$).

As mentioned before, a fixed user profile was used for each domain. Some of them were predefined profiles, and some of them were created by the users (those marked with ‘*’ in the table) using the profile editor of News@hand. In addition, some tasks were done with user profiles containing concepts belonging to all the three domains. They are marked with an ‘A’ in the table.

User	Personalised recommendations		Context-aware recommendations	
	Without expansion	With expansion	With expansion	
	$w_p=1, w_c=0$	$w_p=1, w_c=0$	$w_p=0, w_c=1$	$w_p=0.5, w_c=1$
1	*Q _{1,1}	Q _{2,1}	Q _{3,1}	^A Q _{1,2}
2	Q _{2,2}	*Q _{3,2}	^A Q _{2,1}	Q _{1,2}
3	Q _{3,1}	^A Q _{3,2}	*Q _{1,1}	Q _{2,1}
4	^A Q _{1,1}	Q _{1,2}	Q _{2,2}	*Q _{3,2}
5	Q _{1,2}	*Q _{2,2}	Q _{3,2}	^A Q _{2,1}
6	Q _{2,1}	Q _{3,1}	* ^A Q _{3,2}	Q _{1,1}
7	Q _{3,2}	^A Q _{1,1}	Q _{1,2}	*Q _{2,2}
8	* ^A Q _{2,2}	Q _{1,1}	Q _{2,1}	Q _{3,1}
9	Q _{1,1}	Q _{2,1}	*Q _{3,1}	^A Q _{3,2}
10	Q _{2,2}	Q _{3,2}	^A Q _{1,1}	*Q _{1,2}
11	*Q _{3,1}	^A Q _{2,2}	Q _{1,1}	Q _{2,1}
12	^A Q _{3,1}	*Q _{1,2}	Q _{2,2}	Q _{3,2}
13	Q _{1,2}	Q _{2,2}	Q _{3,2}	* ^A Q _{1,1}
14	*Q _{2,1}	Q _{3,1}	^A Q _{2,2}	Q _{1,1}
15	Q _{3,2}	* ^A Q _{3,1}	Q _{1,2}	Q _{2,2}
16	^A Q _{1,2}	Q _{1,1}	*Q _{2,1}	Q _{3,1}

Table 4. Experiment tasks configurations

Evaluation of personalised recommendations

The objective of the two tasks performed in the first 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.

- Launch the query with personalisation deactivated.
- Rate the top 15 news items. The allowed rating values were: 0 if the item was not relevant to the task goal, 3 if the item was relevant to the task goal, and 5 if the item was relevant to the task goal and the user profile. These ratings, obtained without personalisation and semantic expansion, are considered as our *baseline case*.
- Launch the query with personalisation activated (and semantic expansion enabled/disabled depending on the case).
- Rate again the top 15 items as explained before.

Evaluation of context-aware recommendations

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

- Launch the query with contextualisation deactivated.
- 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.
- Launch the query with the contextualisation activated (semantic expansion enabled, and personalisation enabled/disabled depending on the case).
- Rate again the top 15 items as explained before.

5.4. Results

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

$$P@N = 1/N \cdot \#\{\text{relevant items in the top } N \text{ items}\}$$

Figure 6 shows the average results for the 16 users, taking into account those items evaluated as relevant to the task goal, and also to the user profile. In both cases, the recommendation models outperformed the baseline case, especially for the first 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%.

Although larger scale experiments have to be conducted to claim for statistical significance, the obtained results reinforce conclusions that were previously observed. Personalised recommendations help the user to find relevant news articles, and semantic expansion of user profiles ease the matching between user preferences and item annotations, improving precision values for the top suggested items.

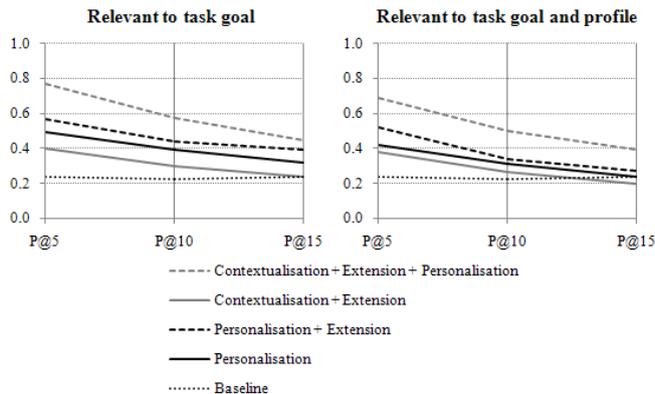


Figure 6. Average precision values for the top 5, 10 and 15 news items, taking into account those items evaluated as relevant to the task goal and the user profile

6. Summary and future work

We have presented News@hand, a news recommender system which makes use of semantic technologies to provide several recommendation services for single and multiple users. In this paper, we describe experiments to preliminary assess their performance once they have been integrated and combined within the system. Specifically, we have evaluated those models that take into account the user's interest profile and the current semantic context.

We are gathering more news contents and user rating information in order to better evaluate the previous and the collaborative recommenders. We are also integrating and testing automatic mechanisms that learn user preferences from semantic contexts and freely-defined social tags.

References

- [1] Alfonseca, E. Moreno-Sandoval, A., Guirao, J. M., Ruiz-Casado, M. (2006). The Wraetlic NLP Suite. *5th Int. Conference on Language Resources and Evaluation (LREC 2006)*.
- [2] Ahn, J., Brusilovsky, P., Grady, J., He, D., Syn, S. Y. (2007). Open User Profiles for Adaptive News Systems: Help or Harm? *16th Int. World Wide Web Conference (WWW 2007)*, pp. 11-20.
- [3] Billsus, D. (1999). A Personal News Agent that Talks, Learns and Explains. *3rd International Conference on Autonomous Agents (Agents 1999)*, 268-275.
- [4] Billsus, D., Pazzani, M. J. (2000). User Modeling for Adaptive News Access. *User Modeling and User-Adapted Interaction*, 10, pp. 147-180.
- [5] Cantador, I., Bellogin, A., Castells, P. (2007). Modelling Ontology-based Multilayered CoI for Hybrid Recommendations. *1st Int. Workshop on Adaptation and Personalisation in Social Systems: Groups, Teams, Communities (SociUM 2007)*.
- [6] Cantador, I., Castells, P. (2006). Multilayered Semantic Social Networks Modelling by Ontology-based User Profiles Clustering: Application to CF. *15th Int. Conf. on Knowledge Engineering and Knowledge Management (EKAW 2006)*, pp. 334-349.
- [7] Cantador, I., Castells, P., Vallet, D. (2006). Enriching Group Profiles with Ontologies for Knowledge-Driven Collaborative Content Retrieval. *1st Int. Workshop on Semantic Technologies in Collaborative Applications (STICA 2006)*, pp. 358-363.
- [8] Cantador, I., Szomszor, M., Alani, H., Fernández, M., Castells, P. (2008). Enriching Ontological User Profiles with Tagging History for Multi-Domain Recommendations. *1st Int. Workshop on Collective Intelligence and the Semantic Web (CISWeb 2008)*.
- [9] Castells, P., Fernández, M., Vallet, D. (2007). An Adaptation of the Vector-Space Model for Ontology-based Information Retrieval. *IEEE TKDE*, 19 (2), pp. 261-272.
- [10] Claypool, M., Gokhale, A., Miranda, T., Murmikov, P., Netes, D., Sartin, M. (1999). Combining Content-Based and Collaborative Filters in an Online Newspaper. *SIGIR'99 Workshop on RS*.
- [11] Das, A., Datar, M., Garg, A., Rajaram, S. (2007). Google News Personalisation: Scalable Online Collaborative Filtering. *16th Int. World Wide Web Conference (WWW 2007)*, pp. 11-20.
- [12] Jones, G. J. F., Queded, D. J., Thomson, K. E. (2000). Personalised Delivery of News Articles from Multiple Sources. In J. Borbinha, T. Baker (Eds.), *ECDL 2000*, pp.340-343.
- [13] Konstan, J. A., Miller, B. N., Maltz, D., Herlocker, J. L., Gordon, L. R., Riedl, J. (1997). GroupLens: Applying CF to Usenet News. *Communications of the ACM*, 40(3), pp. 77-87.
- [14] Lang, K. (1995). NewsWeeder: Learning to filter Netnews. *12th Int. Conference on Machine Learning*, pp: 331-339.
- [15] Maes, P. (1994). Agents that Reduce Work and Information Overload. *Communications of the ACM*, 37(7), pp. 31-40.
- [16] Nadjarbashi-Noghani, M., Zhang, J., Sadat, H., Ghorbani, A.A. (2005). PENS: A Personalised Electronic News System. *3rd Com. Networks and Services Research Conference (CNSR'05)*, pp. 31-38.
- [17] Vallet, D., Castells, P., Fernández, M., Mylonas, P., Avrithis, Y. (2007). Personalised Content Retrieval in Context Using Ontological Knowledge. *IEEE TCSVT*, 17 (3), pp. 336-346.