A RECOMMENDATION SYSTEM FOR THE SEMANTIC WEB

Victor Codina and Luigi Ceccaroni,

Departament de Llenguatges i Sistemes Informàtics (LSI),

Universitat Politècnica de Catalunya (UPC),

Campus Nord, Edif. Omega, C. Jordi Girona, 1-3, 08034 Barcelona, Spain

{vcodina; luigi}@lsi.upc.edu

Abstract Recommendation systems can take advantage of semantic reasoning-capabilities to overcome common limitations of current systems and improve the recommendations' quality. In this paper, we present a personalized-recommendation system, a system that makes use of representations of items and user-profiles based on ontologies in order to provide semantic applications with personalized services. The recommender uses domain ontologies to enhance the personalization: on the one hand, user's interests are modeled in a more effective and accurate way by applying a domain-based inference method; on the other hand, the matching algorithm used by our content-based filtering approach, which provides a measure of the affinity between an item and a user, is enhanced by applying a semantic similarity method. The experimental evaluation on the Netflix movie-dataset demonstrates that the additional knowledge obtained by the semantics-based methods of the recommender contributes to the improvement of recommendation's quality in terms of accuracy.

Keywords: Recommendation systems, Semantic Web, Ontology-based representation, Semantic reasoning, Content-Based filtering, Services Orientation.

1 Introduction

Most common limitations of current recommendation systems are: *cold-start*, *sparsity*, *overspecialization* and *domain-dependency* [4]. Although some particular combination of recommendation techniques can improve the recommendation's quality in some domains, there is not a general solution to overcome these limitations. The use of *semantics* to formally represent data [1] can provide several advantages in the context of personalized recommendation systems, such as the dynamic contextualization of user's interests in specific domains and the guarantee of interoperability of system resources. We think that the next generation of recommenders should focus on how their personalization processes can take advantage of semantics as well as social data to improve their recommendations. In this paper, we describe how the accuracy of recommendation systems is higher when semantically-enhanced methods are applied.

The structure of the paper is as follows: in section 2 we present the state of the art of recommendation systems and semantic recommenders; in section 3 we describe a new domain-independent recommendation system; and in section 4 we present an experimental evaluation of the recommender.

2 Related Work

Different recommendation approaches have been developed using a variety of methods. A detailed review of the traditional approaches based on user and item information, and also a description of the current trend in systems that try to incorporate contextual information to the recommendation process is presented in section 2.3 of Codina [4]. *Semantic* recommendation systems are characterized by the incorporation of semantic knowledge in their processes in order to improve recommendation's quality.

Most of them aim to improve the user-profile representation (*user modeling* stage), employing a concept-based approach and using standard vocabularies and ontology languages like OWL. Two different methods can be distinguished:

- Approaches employing *spreading activation* to maintain user interests and treating the user-profile as a semantic network. The interest scores of a set of concepts are propagated to other related concepts based on pre-computed weights of concepts relations. A news recommender system [3] and a search recommender [8] employ this method.
- Approaches that apply *domain-based inferences*, which consist of making inferences about user's interests based on the hierarchical structure defined by the ontology. The most commonly used is the *upward-propagation*, whose main idea is to assume that the user is interested in a general concept if he is interested in a given percentage of its direct sub-concepts. This kind of mechanisms allows inferring new knowledge about the long-term user's interests and therefore modeling richer user-profiles. *Quickstep* [7], a scientific-paper recommender, and *Travel Support System* [6], a tourism-domain recommender, employ an *upward-propagation* method to complete the user profile.

Other recommenders focus on exploiting semantics to improve the *content adaptation* stage. Most of them make use of *semantic similarity* methods to enhance the performance of a *content-based* approach (CB), although there are also some recommenders using semantics to enhance the user-profile matching of a *collaborative filtering* approach. The only recommender that makes use of semantic reasoning methods in both stages of the personalization process is *AVATAR* [2], a TV recommender that employs *upward-propagation* and *semantic similarity* methods.

3 A Semantic Recommendation System

In this section we present the main components and characteristics of the semantic recommendation system we developed, which makes use of semantics-based methods to enhance both stages of the personalization process.

3.1 Architectural Design

In order to develop a domain-independent recommender, it is necessary to decouple the recommendation engine from the application domains. For this reason, we designed the system as a service provider following the well-known *service oriented architecture* (SOA) paradigm. In **Fig. 1**, the abstract architectural design is represented. Using this decoupled design, each Web-application or domain has to expose a list of items to be used in the personalization process; items has to be semantically annotated using the hierarchically structured concepts of the domain ontology, which is shared with the recommender. Thus, the recommendation engine can work as a personalization service, providing methods to generate personalized recommendations as well as to collect user feedback while users interact with Web-applications. In order to facilitate the reuse of user profiles as well as the authentication process we employ the widely used FOAF vocabulary as the basis of our ontologically extended user profiles, which is compatible with the *OpenID* authentication [http://openid.net/].

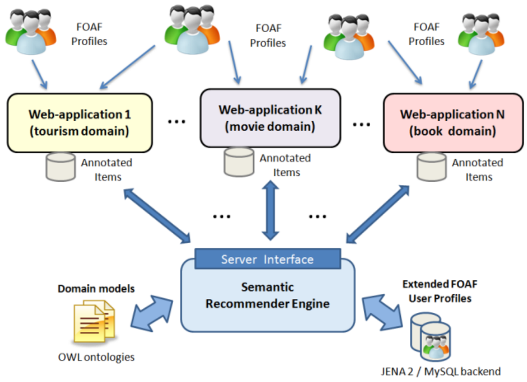

Fig. 1. General architecture design

3.2 Semantic Reasoning Method

Our semantic recommender employs the typical weighted overlay approach, used in ontological user profiles to model user's interests, that consists of mapping collected feedback about semantically annotated items to the corresponding concepts of the domain; the association is done with a weight, which indicates the degree of interest (*DOI_weight*) of the user. In combination with the weight value, we use a measure of how trustworthy is the interest prediction of the particular concept (*DOI_confidence*) to reduce/increase its influence during the recommendation. The recommender takes advantage of this ontological representation in the two stages of the personalization process:

- The user-profile learning algorithm, responsible for expanding and maintaining up-to-date the long-term user's interests, employs a *domain-based inference* method in combination with other relevance feedback methods to populate more quickly the user profile and therefore reduce the typical cold-start problem.
- The filtering algorithm, which follows a CB approach, makes use of a *semantic similarity* method based on the hierarchical structure of the ontology to refine the item-user matching score calculation.

3.2.1 The Domain-Based Inference Method

The domain-based inference method we used is an adaptation of the approach presented in [5] and consists of inferring the degree of interest for a concept using subclass or sibling relations (upward or sideward propagation) when the user is also interested in a minimum percentage (the inference threshold) of direct sub-concepts or sibling concepts. The predicted weight is calculated as the *DOI_weight* average of the sub-concepts or sibling concepts the user is interested in, and the confidence value is based on the percentage of sub-concepts or siblings used in the inference and the average of their respective *DOI_confidence* values.

In **Fig. 2**, we present a graphical example showing how the domain-based inference method works. In a certain moment, the system knows the user is interested in 4 sub-concepts of the *Sport* class (Baseball, Basketball, Football and Tennis). In this case, the proportion of sub-concepts the user is interested in (4 out of 5, i.e., 0.8) is greater than both inference thresholds, therefore both can be applied. Thus, the system infers that the user is interested in *Sport* and *Golf* with the same *DOI_weight* (0.62). The difference between the two types of inference is that the *DOI_confidence* of the sideward-propagation is lower than the one of the upward-propagation (0.5 vs. 0.66).

3.2.2 The Semantic Similarity Method

The basic idea of this method is to measure the relevance of the matching between a particular concept the user is interested in and a concept describing the item. (In **Fig. 3**, two examples are shown, in which the user's interest is the parent of the item concept.) We can distinguish two types of matching:

- The item concept is one of the user's interests, so the matching is perfect and the similarity is maximum (1).
- An ancestor of the item concept (e.g., the direct parent) is one of the user's interests. In this case the similarity is calculated using the following recursive function whose result is always a real number (lower than 1).

  – $SIM_n = SIM_{n-1} - K * SIM_{n-1} * n$ (partial match, n>0)
  – $SIM_0 = 1$ (perfect match, n=0)

  Where:
  - ***n*** is the distance between the item concept and the user's interest (e.g., when it is the direct parent, ***n*** = 1);
  - ***K*** is the factor that marks the rate at which the similarity decreases (the higher ***n***, the higher the decrement). This factor is calculated taking into account the depth of the item concept in the hierarchy and is based on the assumption that semantic differences among upper-level concepts are bigger than those among lower-level concepts.

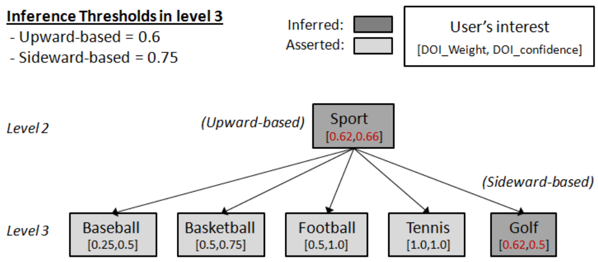

Fig. 2. An example of how new interests are inferred

4 Experimental Evaluation

In this section the undertaken experimental evaluation of the recommender is presented.

The main goal of the experiments is to demonstrate how the recommendation's quality of a CB approach is improved when semantically-enhanced algorithms are employed. We employ the well-known Netflix-prize movie dataset in order to evaluate the recommendation's quality of the recommender in terms of accuracy of rating predictions. The Netflix dataset consists of 480,000 users, 17,700 movies and a total of 100,480,507 user's ratings ranging between 1 and 5. We employ the same predictive-based metric used in the contest, the *root mean square error* (RMSE).

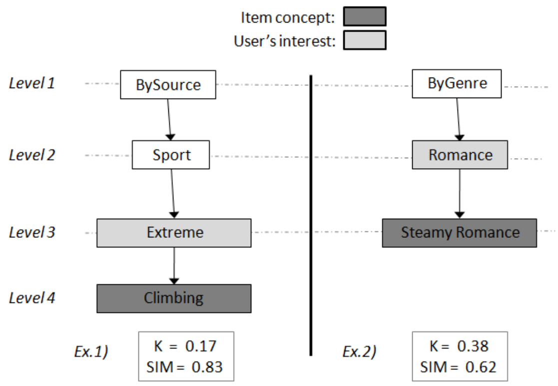

Fig. 3. How the similarity method works

4.2 *Experimental Setup*

To evaluate how the semantically-enhanced algorithms contribute to improve the recommendation's quality in terms of accuracy, we compare the prediction results obtained executing the recommender in three different configurations:

- *CB*. It represents the traditional CB approach; therefore the methods that take advantage of the ontology representation are disabled. In this case, the item-user matching only takes into account the concepts that perfectly match.
- *Sem-CB*. It employs the semantics-based methods presented in section 3 using, as domain ontology and movie indexation, the same taxonomy of three levels of depth used by Netflix and publicly available [http://www.netflix.com/AllGenresList].
- *Sem-CB+*. It employs the semantic-based methods using, as domain ontology, an adaptation of the Netflix taxonomy, with a concepts hierarchy of four levels of depth (see **Fig. 4**). We also changed the indexation for concepts referring to two or more other concepts (i.e., we indexed movies related to Netflix's concept "*Family Dramas*" separately under "*Family*" and "*Drama*") in order to reduce the ontology size.

4.3 *Results*

The error of the predictions generated by the system (see **Table 1**) demonstrates that, when semantics is used, the recommendation's accuracy improves with respect to the *CB* configuration. The accuracy of *Sem-CB+* is not better than *Sem-CB* when the parameters of the algorithms are properly adjusted (see Ex. 3 in **Table 2**). We compare both configurations using the same inference thresholds and the value of the K factor which provides the best accuracy in each case. In the case of *Sem-CB*: K=0.12 when the concept level is 3; K=0.31 when the level is 2. In the case of *Sem-CB+*: K=0.30 when the level is 4; K=0.40 when the level is 3; and K=0.50 when the level is 2. It can be observed that the improvement of accuracy is strongly related with the upward-inference threshold (the higher the number of upward-propagations, the better the results). For example, for Sem-CB+: 1.0443 – 1.0425 – 1.0397.

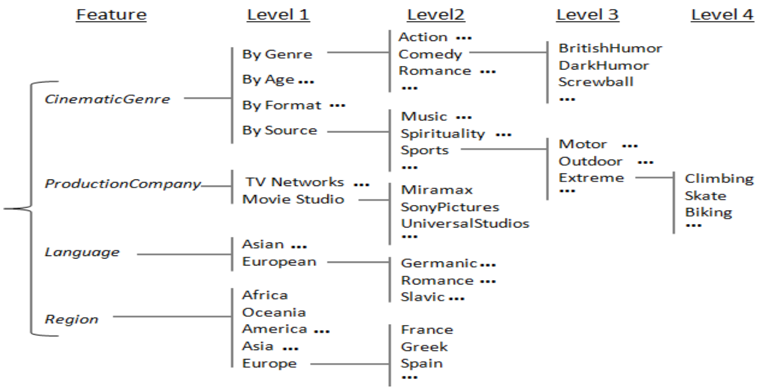

Fig. 4. Partial representation of the adapted movie taxonomy

For comparison, a trivial algorithm that predicts for each movie in the quiz set its average grade from the training data produces an RMSE of 1.0540. Netflix's *Cinematch* algorithm uses "straightforward statistical linear models with a lot of data conditioning". Using only the training data, Cinematch scores an RMSE of 0.9514 on the quiz data, roughly a 10% improvement over the trivial algorithm.

Table 1. Global prediction-error (RMSE) results

Configuration	RMSE
CB	1.0603
Sem-CB	1.0391
Sem-CB+	1.0397

Table 2. Comparison of semantic-based configurations

Execution (*Upward – Sideward*) thresholds		Avg. Upward propagations	Avg. Sideward propagations	RMSE
Ex. 1 (0.60-0.75)	*Sem-CB*	4.32	2.87	1.0482
	Sem-CB+	6.01	3.83	1.0443
Ex. 2 (0.40-0.75)	*Sem-CB*	8.89	3.85	1.0440
	Sem-CB+	9.99	3.89	1.0425
Ex. 3 (0.20-0.85)	*Sem-CB*	13.84	2.88	**1.0391**
	Sem-CB+	17.73	3.30	**1.0397**

5 Conclusions and Future Work

This paper describes how the accuracy of recommendation systems is higher when semantically-enhanced methods are applied. In our approach, we make use of semantics by applying two different methods. A domain-based method makes inferences about user's interests and a taxonomy-based similarity method is used to refine the item-user matching algorithm, improving overall results.

The recommender proposed is domain-independent, is implemented as a Web service, and uses both explicit and implicit feedback-collection methods to obtain information on user's interests. The use of a FOAF-based user-model linked with concepts of domain ontologies allows an easy integration of the recommender into Web-applications in any domain.

As future work we plan to add a collaborative-filtering strategy that makes use of domain semantics to enhance the typical user-profile similarity methods.

6 References

1. Berners-Lee T., J. Hendler, and O. Lassila. 2001. The Semantic Web. A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities. *Scientific American* 284 (5), 34-43.
2. Blanco-Fernández, Y. et al. 2008. A flexible semantic inference methodology to reason about user preferences in knowledge-based recommender systems. *Knowledge-Based Systems* 21 (4), 305-320.
3. Cantador, I., Bellogín, A., and Castells, P. 2008. Ontology-based personalised and context-aware recommendations of news items. *Proc. of IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, 562–565.
4. Codina, V. 2009 Design, development and deployment of an intelligent, personalized recommendation system. Master Thesis. Departament de Llenguatges i Sistemes Informàtics, Universitat Politècnica de Catalunya. 101 pp.
5. Fink, J. and Kobsa, A. 2002. User Modeling for Personalized City Tours. *Artificial Intelligence Review* 18 (1), 33-74.
6. Gawinecki, M., Vetulani, Z., Gordon, M., Paprzycki, M. 2005. Representing users in a travel support system. *Proceedings - 5th International Conference on Intelligent Systems Design and Applications 2005, ISDA '05*, art. no. 1578817, 393-398.
7. Middleton, S.E., De Roure, D. C., and Shadbolt, N.R. 2001. Capturing Knowledge of User Preferences: ontologies on recommender systems. *In Proceedings of the First International Conference on Knowledge Capture (K-CAP 2001), Oct 2001, Victoria, B.C. Canada*.
8. Sieg, A., Mobasher, B., Burke, R. 2007. Ontological user profiles for personalized Web search. *AAAI Workshop - Technical Report WS-07-08*, 84-91.