WebCORE: A Web application for Collaborative Ontology Reuse and Evaluation

Iván Cantador, Miriam Fernández, Pablo Castells
Escuela Politécnica Superior
Universidad Autónoma de Madrid
Campus de Cantoblanco, 28049, Madrid, Spain
{ivan.cantador, miriam.fernandez, pablo.castells}@uam.es

ABSTRACT
In this work we present WebCORE: a Web application for collaborative ontology reuse and evaluation. The system receives an informal description of a specific semantic domain and determines which ontologies from a repository are the most appropriate to describe the given domain. For this task, the environment is divided into three modules. The first component receives the problem description as a set of terms, and allows the user to refine and enlarge it using WordNet. The second module applies multiple automatic criteria to evaluate the ontologies of the repository, and determines which ones fit best the problem description. A ranked list of ontologies is returned for each criterion, and the lists are combined by means of rank fusion techniques. Finally, the third component uses manual user evaluations in order to incorporate a human, collaborative assessment of the ontologies. The new version of the system incorporates several novelties, such as its implementation as a web application; the incorporation of a NLP module to manage the problem definitions; modifications on the automatic ontology retrieval strategies; and a collaborative framework to find potential relevant terms according to previous user queries. Finally, we present some early experiments on ontology retrieval and evaluation, showing the benefits of our system.

Categories and Subject Descriptors
H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval – *information filtering, retrieval models, selection process*.

General Terms
Algorithms, Measurement, Human Factors.

Keywords
Ontology evaluation, ontology reuse, rank fusion, collaborative filtering, WordNet.

1. INTRODUCTION
The Web can be considered as a live entity that grows and evolves fast over time. The amount of content stored and shared on the web is increasing quickly and continuously. The global body of multimedia resources on the Internet is undergoing a significant growth, reaching a presence comparable to that of traditional text contents. The consequences of this enlargement result in well known difficulties and problems, such as finding and properly managing all the existing amount of sparse information.

To overcome these limitations the so-called "Semantic Web" trend has emerged with the aim of helping machines process information, enabling browsers or other software agents to automatically find, share and combine information in consistent ways. At the core of these new technologies, ontologies are envisioned as key elements to represent knowledge that can be understood, used and shared among distributed applications and machines. However, ontological knowledge mining and development are difficult and costly tasks that require major engineering efforts. In this context, ontology reuse becomes an essential need in order to exploit past and current efforts and achievements.

Novel tools have been recently developed, such as as ontology search engines [6] represent an important first step towards automatically assessing and retrieving ontologies which satisfy user queries and requests. However, ontology reuse demands additional efforts to address special needs and requirements from ontology engineers and practitioners. It is necessary to evaluate and measure specific ontology features, such as lexical vocabulary, relations [3], restrictions, consistency, correctness, etc., before making an adequate selection. Some of these features can be measured automatically, but some, like the correctness or the level of formality, require a human judgment to be assessed.

Following this aspiration, the work presented here aims to enhance ontology retrieval and recommendation, combining automatic evaluation techniques with explicit users' opinions and experiences. This work follows a previous approach for Collaborative Ontology Reuse and Evaluation over controlled repositories, named CORE [2]. For the work reported in this paper, the tool has been enhanced and adapted to the Web. Novel technologies, such as AJAX,[1] have been incorporated to the system for the design and implementation of the user interface. It has also been modified and improved to overcome previous limitations, such as handling large numbers of ontologies. The collaborative capabilities have also been extended within two different frameworks. Firstly, during the problem definition phase, the system helps users to express their needs and requirements by showing other problem descriptions previously given by different users. Secondly, during the ontology retrieval phase, the system helps users to enhance the automatic system recommendations by using other user evaluations and comments.

The rest of the paper has been organized as follows. The system architecture is described in Section 2 and some conclusions are reported in section 3.

[1] AJAX, http://en.wikipedia.org/wiki/Ajax_(programming)

2. SYSTEM ARCHITECTURE

As mentioned before, WebCORE is a web application for Collaborative Ontology Reuse and Evaluation. A user logins into the system via a web browser, and, thanks to AJAX technology and the Google Web Toolkit[2], dynamically describes a problem domain, searches for ontologies related to this domain, obtains relevant ontologies ranked by several lexical, taxonomic and collaborative criteria, and optionally evaluates by himself those ontologies that he likes or dislikes most.

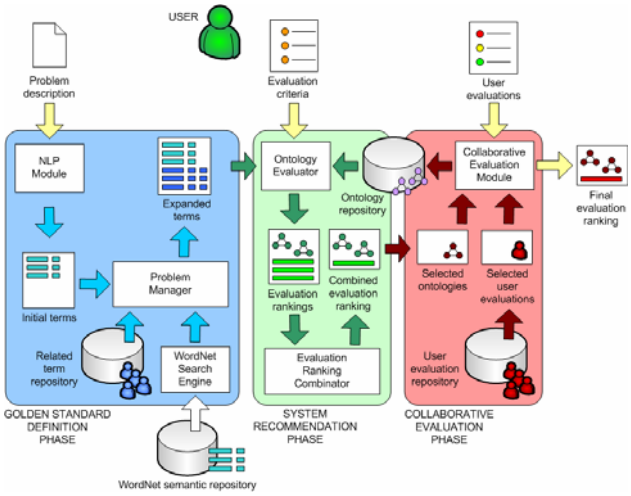

Figure 1. WebCORE architecture

In this section, we describe the server-side architecture of WebCORE. Figure 1 shows an overview of the system. We distinguish three different modules. The first one, the left module, receives the problem description (Golden Standard) as a full text or as a set of initial terms. In the first case, the system uses a NLP module to obtain the most relevant terms of the given text. The initial set of terms can also be modified and extended by the user using WordNet [4]. The second one, represented in the centre of the figure, allows the user to select a set of ontology evaluation techniques provided by the system to recover the ontologies closest to the given Golden Standard. Finally, the third one, on the right of the figure, is a collaborative module that re-ranks the list of recovered ontologies, taking into consideration previous feedback and evaluations of the users.

2.1 Golden Standard Definition

The first phase of our ontology recommender system is the Golden Standard definition. The user describes a domain of interest specifying a set of relevant terms that will be searched through the concepts (classes or instances) of the ontologies stored in the system. As an improvement of [2], in WebCORE we have added a new collaborative component that continuously offers to the user a ranked list with the terms that have been used in those previous problem descriptions in which a given term appears.

Figure 2 shows the interface of the Golden Standard Definition phase. In the left side of the screen, the current list of root terms is shown. The user can manually insert new root terms to this list giving their lexical entries and selecting their parts of speech. The correctness of these new insertions is controlled by verifying that all the considered lexical entries belong to the WordNet repository. Adding new terms, the final Golden Standard definition is immediately updated: the final list of (root and expanded) terms that represent the domain of the problem is shown in the bottom of the figure. The user can also make term expansion using WordNet. He selects one of the terms from the Golden Standard definition and the system shows him all its meanings contained in WordNet (top of the figure). After he has chosen one of them, the system presents him three different lists with the synonyms, hyponyms and hypernyms of the term. The user can then select one or more elements of these lists and add them to the expanded term list. For each expansion, the depth of the new term is increased by one unit. This will be used later to

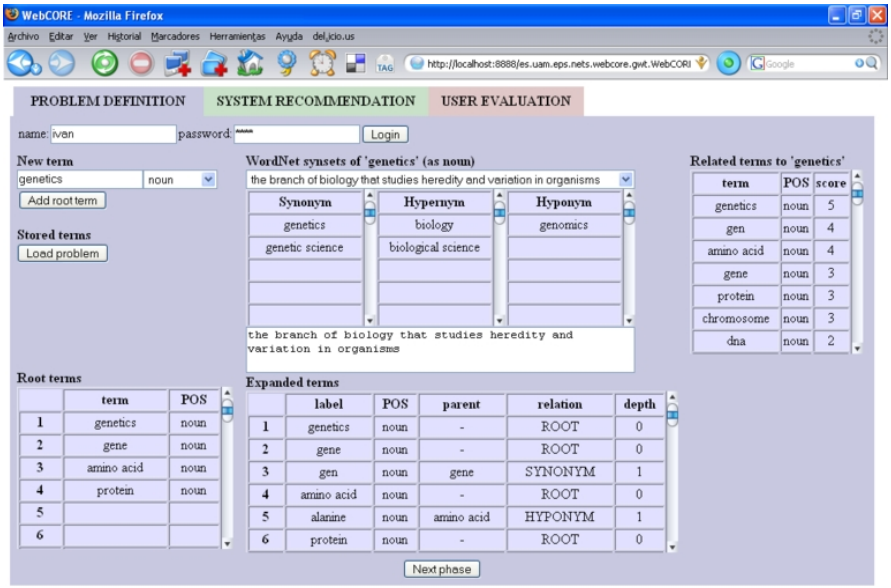

Figure 2. WebCORE problem definition phase

[2] Google Web Toolkit, http://code.google.com/webtoolkit/

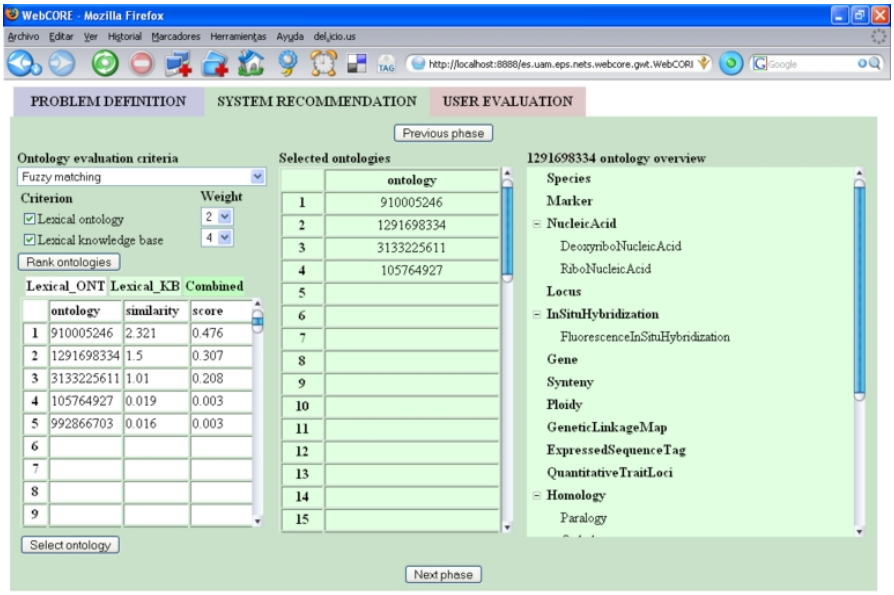

Figure 3. WebCORE system recommendation phase

measure the importance of the term within the Golden Standard: the greater the depth of the derived term with respect to its root term, the less its relevance will be.

In the problem definition phase a collaborative component has been added to the system (right side of Figure 2). This component reads the term currently selected by the user, and searches for all the stored problem definitions that contain it. For each of these problem definitions, the rest of their terms and the number of problems in which they appear are retrieved and shown in the web browser. With this simple strategy the user is suggested the most popular terms, fact that could help him to better describe the domain in which he is interested in.

2.2 Automatic Ontology Recommendation

Once the user has selected the most appropriate set of terms to describe the problem domain, the tool performs the processes of ontology retrieval and ranking. Our approach to ontology retrieval can be seen as an evolution of classic keyword-based retrieval techniques [5], where textual documents are replaced by ontologies.

The queries supported by our model are expressed using the terms selected during the Golden Standard definition phase. In classic keyword-based vector-space models for information retrieval [5], the query keywords are assigned a weight that represents the importance of the concept in the information need expressed by the query. Analogously, in our model, the terms included in the Golden Standard are weighted, using the depth measure, to indicate the relative interest of the user for each of the terms to be explicitly mentioned in the ontologies.

To carry out the process of ontology retrieval, the approach is focused on the lexical level, retrieving those ontologies that contain a subset of the terms expressed by the user during the Golden Standard definition. To compute the matching, two different options are available within the tool: search for exact matches or search for matches based on the Levenshtein distance between two terms. The tool also offers two different search spaces, the ontologies and the corresponding knowledge bases. The retrieval process returns a set of ontologies that satisfy user requirements. This set of ontologies is ranked by an algorithm that adapts the vector-space model principles to a vector representation of ontology concepts.

Figure 3 shows the system recommendation interface. At the right side the user can select the matching methodology (fuzzy or exact), the search spaces (ontology entities and knowledge base entities), and the weight or importance given to each of the previously selected search spaces. In the right part the user can visualize the ontology and navigate across it. Finally, the middle of the interface presents the list of ontologies selected for the user to be evaluated during the collaborative evaluation phase.

2.3 Collaborative Ontology Evaluation

The third and last phase of the system is compound of a novel ontology recommendation algorithm that exploits the advantages of Collaborative Filtering [1], exploring the manual evaluations stored in the system to rank the set of ontologies that best fulfils the user's interests.

In WebCORE, user evaluations are represented as a set of five different criteria and their respective values, manually determined by the users who made the evaluations.

- **Correctness**: specifies whether the information stored in the ontology is true, independently of the domain of interest.
- **Readability**: indicates the non-ambiguous interpretation of the meaning of the concept names.
- **Flexibility**: points out the adaptability or capability of the ontology to change.
- **Level of formality**: highly informal, semi-informal, semi-formal, rigorously-formal.
- **Type of model**: upper-level (for ontologies describing general, domain-independent concepts), core-ontologies (for ontologies that contain the most important concepts on a specific domain), domain-ontologies (for ontologies that

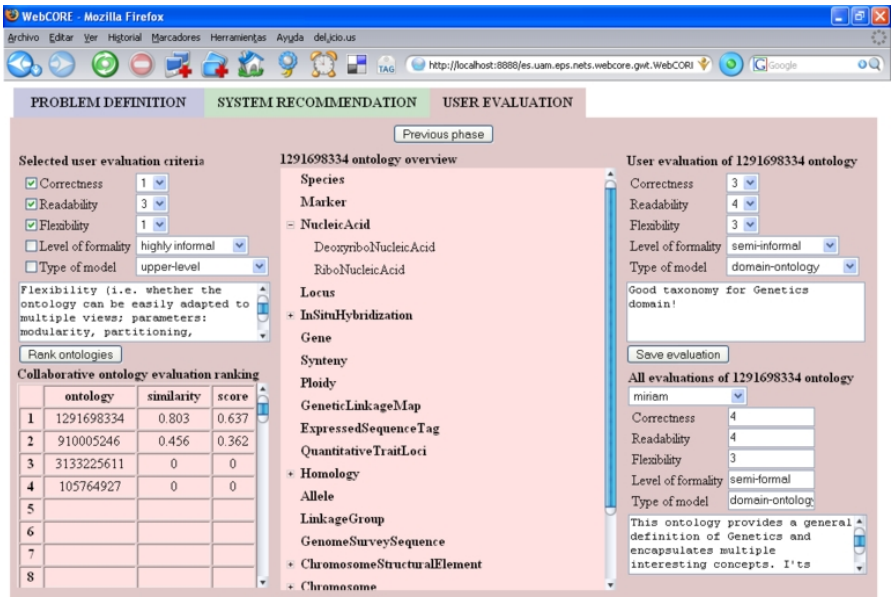

Figure 4. WebCORE user evaluation phase

broadly describe a domain), task-ontologies (for ontologies focused on generic types of tasks or activities) and application-ontologies (for ontologies describing a domain in an application-dependent manner).

The above criteria can have discrete numeric or non-numeric values. The user's interests are expressed like a subset of these criteria, and their respective values, meaning thresholds or restrictions to be satisfied by user evaluations. Thus, a numeric criterion will be satisfied if an evaluation value is equal or greater than that expressed by its interest threshold, while a non-numeric criterion will be satisfied only when the evaluation is exactly the given threshold (i.e. in a Boolean or yes/no manner).

According to both types of user evaluation and interest criteria, numeric and Boolean, the recommendation algorithm will measure the degree in which each user restriction is satisfied by the evaluations, and will recommend a ranked ontology list according to similarity measures between the thresholds and the collaborative evaluations. To create the final ranked ontology list the recommender module follows two phases. In the first one it calculates the similarity degrees between all the user evaluations and the specified user interest criteria thresholds. In the second one it combines the similarity measures of the evaluations, generating the overall rankings of the ontologies.

Figure 4 shows all the previous definitions and ideas, locating them in the graphical interface of the system. On the left side of the screen, the user introduces the thresholds for the recommendations and obtains the final collaborative ontology ranking. On the right side, the user adds new evaluations for the ontologies and checks evaluations given by the rest of the users.

3. CONCLUSIONS

In this paper, a web application for ontology evaluation and reuse has been presented. The novel aspects of our proposal include the use of WordNet to help users to define the Golden Standard; a new ontology retrieval technique based on traditional Information Retrieval models; rank fusion techniques to combine different ontology evaluation measures; and two collaborative modules: one that suggests the most popular terms for a given domain, and one that recommends lists of ontologies with a multi-criteria strategy that takes into account user opinions about ontology features that can only be assessed by humans.

4. ACKNOWLEDGMENTS

This research was supported by the Spanish Ministry of Science and Education (TIN2005-06885 and FPU program).

5. REFERENCES

[1] Adomavicius, G., and Tuzhilin, A.: *Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions*. IEEE Transactions on Knowledge and Data Engineering 17(6): 734-749, 2005.

[2] Fernández, M., Cantador, I., and Castells, P. *CORE: A Tool for Collaborative Ontology Reuse and Evaluation*. Proceedings of the 4th Int. Workshop on Evaluation of Ontologies for the Web (EON'06), at the 15th Int. World Wide Web Conference (WWW'06). Edinburgh, UK, 2006.

[3] Maedche, A., and Staab, S.: *Measuring similarity between ontologies*. Proceedings of the 13th European Conference on Knowledge Acquisition and Management (EKAW 2002). Madrid, Spain, 2002.

[4] Miller, G. A.: *WordNet: A lexical database for English. New horizons in commercial and industrial Artificial Intelligence*. Communications of the Association for Computing Machinery, 38(11): 39-41, 1995.

[5] Salton, G., and McGill, M.: *Introduction to Modern Information Retrieval*. McGraw-Hill, New York, 1983.

[6] *Swoogle - Semantic Web Search Engine*.

http://swoogle.umbc.edu