

Cupboard – A Place to Expose your Ontologies to Applications and the Community*

Mathieu d’Aquin¹ and Holger Lewen²

¹ Knowledge Media Institute (KMi), The Open University, United Kingdom
m.daquin@open.ac.uk

² Institute AIFB, Universität Karlsruhe (TH), Germany
lewen@aifb.uni-karlsruhe.de

Abstract. In this demo, we present the Cupboard system for ontology publishing, sharing and reuse. This system is intended to support both ontology engineers and ontology users/practitioners. For the developers of ontologies, it offers a complete infrastructure to host their ontologies in online ontology spaces, providing mechanisms to describe, manage and effectively exploit these ontologies (through APIs). Furthermore, these ontologies are then exposed to the community, providing users with a complete, friendly environment to find, assess and reuse ontologies.

1 Introduction

In the recent years, there has been an ever-increasing amount of ontologies indexed by ontology search engines such as Watson³. Yet, even with so many ontologies now available to reuse, the tasks of finding, assessing and effectively making use of existing ontologies remain difficult. We argue that the main factor inhibiting ontology reuse relates to the lack of support for ontology practitioners to find, assess and exploit existing ontologies, and incidentally to the lack of support for ontology engineers to effectively deploy and expose their ontologies for use and reuse.

Ontology search engines provide a partial answer to this problem, as they support users in finding, selecting and, in some cases, using ontologies (see e.g., [1]). However, they still lack mechanisms for assessing ontologies, for providing rich ontology metadata, for providing alignments between ontologies, etc. In other terms, they only focus on providing the ontologies, without leaving ontology publishers and users any control or any possibility of interaction among themselves and with the ontologies. In addition, there have been several initiatives coming up with ontology repositories lately (see e.g., [2, 3]). However, these systems are generally limited in their scope (e.g., considering only a particular domain or a particular type of ontologies) and leave only little control to the user over the ontology publication process.

* This work was partially supported by the European Commission in the IST project NeOn (IST-2006-027595)

³ <http://watson.kmi.open.ac.uk>

Apart from the Watson ontology search engine already mentioned, the NeOn project⁴ is developing a number of technologies that, if put together and integrated in a broader system, can provide a more comprehensive answer to the aforementioned problem. More specifically, Oyster [4] is a peer-to-peer ontology sharing system which relies on rich metadata for ontologies, using the OMV format [5]. The alignment server [6] is an open system dedicated to the creation, management and evaluation of ontology alignments. In addition, a *topic specific open rating system* (TS-ORS) [7] for ontologies is being developed to allow users to review different properties of ontologies and to express trust in reviewers.

In this demo, we present Cupboard, an online system that relies on these components and integrates them to provide a complete and friendly environment for ontology publishing, sharing and reuse. One of the most important questions addressed in this development concerns the incentives provided for an ontology engineer to share his/her ontology and invest the time to properly add it to a repository. For many users, setting up their own server infrastructure, generating search indexes and setting up SPARQL-endpoints was quite a cumbersome task. Cupboard does not only allow users to add their ontologies in a personal space, but actually indexes them (using the Watson engine), provides mechanism to link them (using the Alignment Server), hosts them and exposes them through APIs and SPARQL, thus offering the infrastructure for engineers to actually deploy their ontologies with minimal efforts. Moreover, Cupboard is designed to be a community tool, helping ontology users and practitioners (including ontology developers) in finding and reusing ontologies, through the use of rich ontology metadata (thanks to Oyster and OMV) and to advanced ontology review mechanisms (using the TS-ORS system).

2 The Cupboard System

Cupboard is different from classical ontology repository systems as it does not provide one single space where ontologies are exposed. On the contrary, each user can create his/her own *ontology space*, containing the ontologies he/she has selected. In a sense, Cupboard can be seen more as a system to host ontology repositories, than as an ontology repository itself.

2.1 Overview

The main advantage in relying on this notion of ontology spaces is that each ontology space virtually implements a complete infrastructure for building semantic applications. Indeed, each uploaded ontology is automatically indexed using the Watson engine. Watson is designed to be an ontology search engine, but focuses on providing the necessary access mechanisms (through services, APIs, SPARQL, etc.) to enable semantic applications to dynamically exploit ontologies made available online (see [1]). In Cupboard, each ontology space becomes like a little Watson, making it possible for developers to easily build applications exploiting the exposed ontologies.

⁴ <http://neon-project.org>

However, for ontology and application developers, Cupboard’s ontology spaces should do more than listing ontologies. To be used jointly in applications, ontologies have to be related with each other. Using the Alignment Server [6] as a core component, Cupboard allows users to populate their ontology spaces not only with ontologies, but also with alignments (mappings). Using the *Alignment Format* [6], alignments can be uploaded for a given (pair of) ontology(ies). The Alignment server then offers to Cupboard the necessary features to support the management, evaluation, and even production of alignments.

Uploading an ontology into Cupboard is not only a way to build applications, but also a way to share this ontology with the community for reuse. It is therefore important to provide efficient ways for users to find, assess and select ontologies in Cupboard’s ontology spaces. The system naturally inherits the powerful search mechanisms implemented in the Watson engine. In addition, information for example on the provenance of the ontology or the methodology used to build it are essential for ontology selection. Cupboard relies on the Oyster system [4] to provide users with the possibility to enter rich metadata about ontologies (in the OMV format [5]) and to manage these metadata.

Finally, one of the most important issues hampering reuse concerns the evaluation of ontologies. In Cupboard, we intend to exploit the community feature of the system, so that users can provide precise reviews on ontologies (reviewing separately different characteristics, like reusability, complexity, coverage, etc.) The TS-ORS System [7], used to handle and manage these reviews, implements a mechanism to rank ontologies according to these evaluations. However, reviews are subjective evaluations, as users might have different views on ontologies. When searching for an ontology to reuse, not only the review information should be exploited to support the user, but also the information about how much he/she would trust a particular review or a particular reviewer. Such trust information is obtained through the classical “Did you find this review useful?” question, and handled by the TS-ORS system to compute trust information used to personalize ontology ranking.

Figure 1 shows different parts of the Web interface of Cupboard, in particular displaying an ontology space, uploading an ontology, inspecting and searching the content of a particular ontology, reviewing it, etc.

2.2 Architecture

While the different technologies employed and combined in the system have been described and implemented before, Cupboard provides homogeneous (user and programming) interfaces to these components and features. The main challenges during integration were defining the interfaces between components, optimizing the performance and generating an easy-to-use Web interface. However, this has been made easier by the fact that the integrated systems are designed to work as independent, remote components, resulting in a naturally distributed architecture for Cupboard.

As depicted in Figure 2, the central component of Cupboard is a specifically made component, called the *Cupboard core*, which is in charge of orchestrating the interaction between all the other components, external applications and

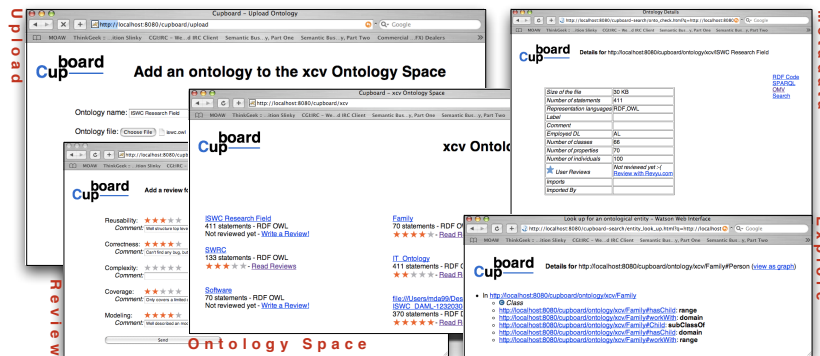


Fig. 1. Screenshots of some components of the Cupboard Web interface.

users. It heavily relies on the Watson engine to provide basic functionalities such as storing the ontologies, indexing them for search, validating them and exposing them to applications. The TS-ORS open rating system, the Oyster ontology metadata sharing system and the Alignment Server each provide specific functionalities to the system, that are federated and exposed in an homogeneous way to the Cupboard Web interface and APIs, through the Cupboard core.

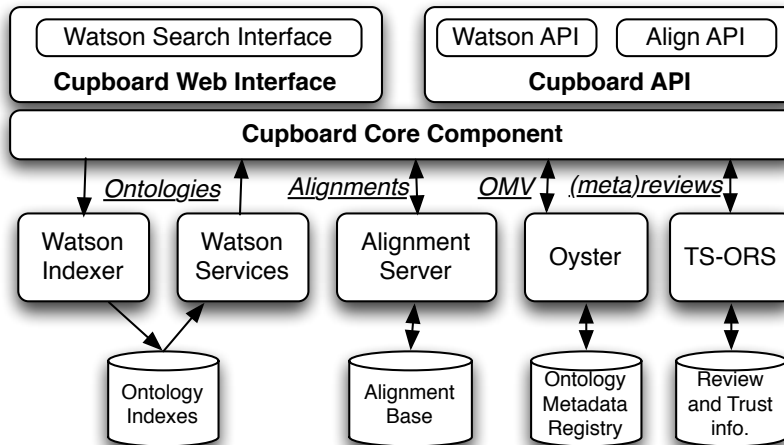


Fig. 2. Cupboard's architecture.

3 Demonstration and Availability

In the demonstration, users will learn how to use Cupboard, to upload, expose and explore their ontologies, as well as to find, select, reuse and assess other users' ontologies. They will learn about the benefits of using Cupboard through a variety of concrete, real-life scenarios, supporting different user groups (ontology engineers, application developers, ontology publishers, etc.) in tackling different tasks (ontology sharing, reuse, application, evaluation, etc.)

At the time of writing, Cupboard is being tested by a selected group of users. It will be made openly available (at <http://cupboard.open.ac.uk>) by June 2009, allowing users to register and create their own ontology spaces. At a later stage, Cupboard will also be distributed as an open source system, to allow organizations to deploy their own Cupboard servers, in case for example intellectual property restrictions prevent them from sharing their ontologies online.

4 Conclusion

In this document, we presented Cupboard, an online ontology publishing and sharing environment that addresses two major problems faced by the Semantic Web community: the deployment of inter-linked ontologies for exploitation in applications, often requiring heavy infrastructures, and the reuse of existing ontologies where, besides search mechanisms, rich metadata and proper assessment of the ontologies are required. To tackle these problems, we propose a complete system, based on a number of well-established technologies, allowing ontology engineers to deploy their ontologies, providing the necessary infrastructures to support their exploitation, and ontology users in reusing available knowledge, providing essential, community-based functionalities to facilitate the search, selection and exploitation of the available ontologies. To the best of our knowledge, Cupboard is the first system to put together all these functionalities to create an essential infrastructure component for Semantic Web developers and more generally, a useful, shared and open environment for the ontology community.

References

1. d'Aquin, M., Motta, E., Sabou, M., Angeletou, S., Gridinoc, L., Lopez, V., Guidi, D.: Toward a new generation of semantic web applications. *Intelligent Systems, IEEE* **23**(3) (2008) 20–28
2. Musen, M., Shah, N., Noy, N., Dai, B., Dorf, M., Griffith, N., Buntrock, J., Jonquet, C., Montegut, M., Rubin, D.: BioPortal: Ontologies and Data Resources with the Click of a Mouse. In: *Proc. of the Annual Symposium proceedings/AMIA Symposium. AMIA Symposium, AMIA Annu Symp Proc* (2008) 1223
3. Buitelaar, P., Eigner, T., Declerck, T.: OntoSelect: A Dynamic Ontology Library with Support for Ontology Selection. In: *Proc. of the Demo Session at the International Semantic Web Conference. Hiroshima, Japan.* (2004)
4. Palma, R., Haase, P.: Oyster - sharing and re-using ontologies in a peer-to-peer community. In: *Proc. of International Semantic Web Conference.* (2005) 1059–1062
5. Hartmann, J., Palma, R., Sure, Y., del Carmen Suárez-Figueroa, M., Haase, P., Gómez-Pérez, A., Studer, R.: Ontology metadata vocabulary and applications. In: *OTM Workshops. Volume 3762 of LNCS., Springer* (2005) 906–915
6. Euzenat, J.: An API for ontology alignment. In McIlraith, S.A., Plexousakis, D., van Harmelen, F., eds.: *International Semantic Web Conference. Volume 3298 of Lecture Notes in Computer Science., Springer* (2004) 698–712
7. Lewen, H., Supekar, K., Noy, N.F., Musen, M.A.: Topic-Specific Trust and Open Rating Systems: An Approach for Ontology Evaluation. In: *Proc. of the 4th International Workshop on Evaluation of Ontologies for the Web (EON2006) at the 15th International World Wide Web Conference (WWW 2006).* (2006)