

Open Ontology Repository

Kenneth Baclawski
College of Computer and Information Science
Northeastern University
Boston, Massachusetts 02115 USA
+1.617.373.4631
kenb@ccs.neu.edu

Todd Schneider
Raytheon
7700 Arlington Blvd.
Fairfax, VA, 22042 USA
+1.703.204.6463
todd.schneider@raytheon.com

ABSTRACT

The Open Ontology Repository is an open source effort to develop infrastructure for ontologies that is federated, robust and secure. This article describes the purpose, requirements and goals of this initiative.

Categories and Subject Descriptors

H.3 [Information Storage and Retrieval]
; H.3.5 [On-line Information Services]

General Terms

Management, Design, Languages, Theory, Legal Aspects

Keywords

Ontology, repository, standard, open, ontology-mapping, community

1. INTRODUCTION

Among the distinguishing characteristics of the semantic web are the use of machines to consume and interpret content, in addition to easing human usage. The use of knowledge representation techniques and mechanisms that facilitate this (e.g. OWL, Common Logic) creates a strong transitive dependency among the documents and content relying on these artifacts (via direct or nested imports). Changes in any of these artifacts can cause the resulting import closure(s) to be inconsistent and/or change the meaning and/or change the computational characteristics.

In order to confidently and consistently use content from the semantic web there must be careful selection and precise reference. Moreover, there must be temporal stability of any references and semantic dependencies of the content. The Open Ontology Repository (OOR) will meet this need.

2. OOR INITIATIVE

The Open Ontology Repository (OOR) [1], an Ontolog Forum community initiative, aims to promote the global use and sharing of ontologies by providing infrastructure for consistency and temporal stability of ontologies for the semantic web. The charter of the Open Ontology Repository includes, (1) Establishing a hosted registry-repository; (2) Enabling and facilitating open, federated, collaborative ontology repositories, and (3) Establishing best practices for expressing interoperable ontology and taxonomy work in registry-repositories. A premise of the OOR is that the

capabilities proposed for the OOR, will facilitate and foster the growth, reuse, and general usage of ontologies and semantic technologies entailed by them. To quote from the Ontology Summit 2008 communiqué,

The purpose of an Open Ontology Repository is to provide an architecture and an infrastructure that supports a) the creation, sharing, searching, and management of ontologies, and b) linkage to database and XML Schema structured data and documents. Complementary goals include fostering the ontology community, the identification and promotion of best practices, and the provision of services relevant to ontologies and instance stores. Examples of anticipated services include automated semantic interpretation of content expressed in knowledge representation languages, the creation and maintenance of mappings among disparate ontologies and content, and inference over this content. We believe that the Open Ontology Repository will ultimately support a broad range of semantic services and applications of interest to enterprises and communities.

This effort was initially promoted by Michael Dean, Mark Musen, Leo Obrst and Peter Yim in January of 2008. There were regular community discussions leading to a 2-day Ontology Summit in April, 2008 [2]. The resulting communiqué "... represents the joint position of those who were engaged in the year's summit discourse on an Open Ontology Repository and of those who endorse [it]." As a community initiative many people from varying institutions are involved. The current list of contributors can be found at [3].

To meet the goals and communiqué intent the OOR community has been engaged in a range of activities including: (1) Developing rationale and motivation, (2) Collecting user needs and requirements, (3) Developing an open architecture, (4) Initiating design and implementation, (5) Identifying research challenges, and (6) Seeking funding. This note focuses on (2) to (5).

3. OOR CREATION

The OOR community understands the need to use a structured development approach in order to meet the needs of the community, gain consensus and acceptance, clearly identify the problems to be solved, and facilitate a consistent development. As with any community effort there were many

contributors and refinement of the contributions. The current set of use cases can be found in [4] and their development is on going.

3.1 Requirements

The current set of requirements can be found at [5]. The following is a subset providing a good understanding of the OOR.

- The repository architecture shall be scalable.
- The repository shall be distributed.
- The specification of the repository shall be sufficiently detailed and platform independent to allow multiple implementations.
- The repository shall be capable of supporting ontologies in languages that have reasoners [supporting inferencing].
- The repository architecture shall support distributed repositories.
- The repository architecture shall not require a hierarchical structure.

The practical outcome of the requirements should result in a repository that provides

- A well-maintained persistent store (with high availability and performance) where ontological work can be stored, shared and accessed consistently;
- Mechanisms for registering and “governing” ontologies, with provenance and versioning, made available (logically) in one place so that they can be browsed, discovered, queried, analyzed, validated and reused;
- Services across disparate ontological artifacts supporting cross-domain interoperability, mapping, application and inferencing; and
- Registration of semantic services to support peer OORs

3.2 Architecture

In order to meet the goals of the OOR, especially common adoption, an open and well documented architecture is required to allow multiple communities and organizations to participate in the OOR and to produce standard OOR functionalities and behaviors. The following architectural principles represent a consensus among the contributors.

- **Decoupling of responsibilities** – To support multiple knowledge representations/languages the repository will not be tightly coupled with the content.
- **Implementation/Platform independence** – To support acceptance, multiple instances, and evolution, no particular implementation or platform dependence can be allowed.
- **Ontologically driven** – To allow for evolution of the OOR and reduce overall development costs, the use of an ontologically based development environment is sought.

3.3 Design and Development

The OOR community understands that the development of the OOR will take place in fits and starts in order to gain ac-

ceptance and properly understand and document the problem space and solutions. An initial prototype design was developed by Mike Dean [6]. The next large step was made when the National Center for Biomedical Ontology (NCBO) [7] contributed their BioPortal, a repository for accessing and sharing ontologies used by biomedical communities, to the OOR for experimentation and testing [8]. The NCBO is actively involved in the OOR initiative with the expectation that results of OOR experimentation will be incorporated into future releases of BioPortal.

4. COMMUNITY SUPPORT

The OOR community anticipates that other groups and organizations will create the actual components to support the OOR functionality for particular languages. One example is Common Logic. Michael Grüninger and his students at the University of Toronto (Ontario, Canada) are developing components that will support the OOR. Another is web service interfaces being developed by Ken Baclawski and his students at Northeastern University. The OOR community welcomes and encourages other groups and organizations, in particular faculty and students, to participate in the creation of the OOR.

5. RESEARCH CHALLENGES

The following is a short list of some of the challenges facing the creation of the OOR as envisioned.

- Ontology metadata
 - Ontology Summit 2007 Dimensions
 - Ontology Metadata Vocabulary
 - eXtended Metadata Registry
- Interface ontologies
 - Internal APIs for core modules and plug-ins
 - External APIs, especially web services
 - Federation APIs, among OORs
- Best practices
 - Policies and procedures
 - Provenance to enable trust

6. ACKNOWLEDGMENTS

Our thanks to the Ontolog Forum and the work of Peter Yim in fostering an open and collaborative environment and providing infrastructure resources to aid the OOR development.

7. REFERENCES

- [1] Open Ontology Repository <http://ontolog.cim3.net/cgi-bin/wiki.pl?OpenOntologyRepository>.
- [2] Ontology Summit 2008 <http://ontolog.cim3.net/cgi-bin/wiki.pl?OntologySummit2008>.
- [3] OOR Contributors <http://ontolog.cim3.net/cgi-bin/wiki.pl?OpenOntologyRepository#nid17Y>.
- [4] OOR Use Cases http://ontolog.cim3.net/cgi-bin/wiki.pl?OpenOntologyRepository_UseCase.
- [5] OOR Requirements http://ontolog.cim3.net/cgi-bin/wiki.pl?OpenOntologyRepository_Requirements.
- [6] <http://projects.semwebcentral.org/projects/oor>.
- [7] NCBO Website <http://bioontology.org>.
- [8] OOR Instance <http://oor-01.cim3.net>.