

Prospects and Possibilities for Ontology Evaluation: The View from NCOR

Leo Obrst

The MITRE Corporation
7515 Colshire Drive
McLean, VA 22102-7508, USA
703-983-6770

lobrst@mitre.org

Todd Hughes

Lockheed Martin Advanced
Technology Laboratories
3 Executive Campus
Cherry Hill, NJ 08002
856-792-9756

thughes@atl.lmco.com

Steve Ray

The US Department of Commerce
National Institute of Standards and
Technology
100 Bureau Dr.
Gaithersburg, MD 20899
301-975-3524

ray@nist.gov

ABSTRACT

In this position paper, we briefly describe the perspective of the US National Center for Ontological Research (NCOR, <http://ncor.us>) on ontology evaluation. NCOR's inauguration was recently held (October 2005), and at that time goals were identified and committees formed to pursue those goals, including the Ontology Evaluation Committee. This committee is charged with developing a plan for the evaluation of ontologies that is designed to transform ontological engineering into a true scientific and engineering discipline. This paper discusses some issues on ontology evaluation, including the relevant questions to ask, and suggests some approaches.

Categories and Subject Descriptors

I.2.4 [Knowledge Representation Formalisms and Methods (F.4.1)]: Predicate logic, Representations (procedural and rule-based)

I.2.3 [Deduction and Theorem Proving (F.4.1)]: Inference engines, Deduction (e.g., natural, rule-based)

I.2.0 [Philosophical Foundations]

General Terms

Experimentation, Standardization, Languages, Theory, Verification.

Keywords

Ontology evaluation, upper ontologies, domain ontologies, ontology representation, ontology reasoning, formal ontology.

1. INTRODUCTION

This position paper describes the perspective of the US National Center for Ontological Research (NCOR) on ontology evaluation. NCOR's inauguration was recently held [1]. At that time goals were identified and committees formed to pursue those goals, including the Ontology Evaluation Committee. This committee is charged with developing a plan for the evaluation of ontologies, to move ontological engineering into a true scientific and engineering discipline.

2. CURRENT STATE OF PRACTICE

Currently, the ontology community resembles the fabled blind men and the elephant: describing the ontology elephant from their own individual perspectives.

Figure 1 displays a diversity of opinion, from how ontology languages should be constituted, to different notions of upper ontologies, to whether ontologies observe the Open World or the Closed World assumption, etc. This state holds because the ontology community is not homogeneous, but represents multiple communities that are converging on the use of semantic technologies, each with a distinct perspective about what an ontology is. Some of the communities are:

- Logicians, formal ontologists, formal semanticists, some computer scientists
- Librarians, information scientists
- Object-oriented software engineers
- Classical artificial intelligence knowledge engineers
- Database theorists and practitioners
- The World Wide Web community
- Enterprise architects and others involved in Service Oriented Architecture and Web service development
- Business and government analysts
- Domain experts

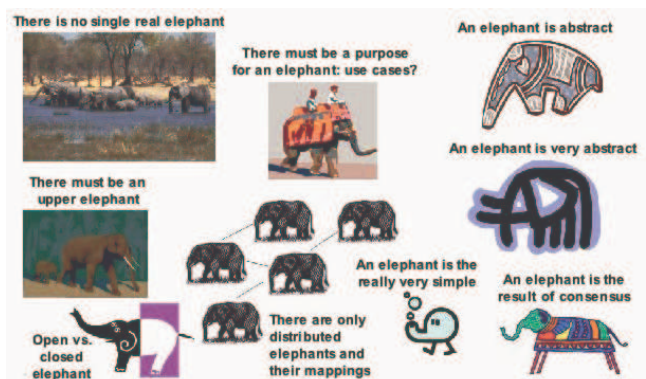


Figure 1. Ontology Elephants

Because the community is heterogeneous, key ontological distinctions are typically glossed over: term vs. concept, label vs. model, machine vs. human interpretability, syntax vs. semantics-pragmatics (sense, reference, discourse, speech acts), philosophical stance (realist, idealist, nominalist), etc.

3. NCOR'S GOALS

NCOR's primary purpose is to help evolve the state of the art and practice toward creating a science and engineering discipline to improve the quality of ontologies used in information technology. To achieve this purpose, NCOR has the following goals:

1. Advance the quality of knowledge/semantic representation languages, ontology content, and runtime reasoning methods.
2. Create procedures, processes, methods to help define, adjudicate, and ensure quality of knowledge.
3. Facilitate the education of communities and promote best practices for ontology development.
4. Promote the best standards related to ontologies, and facilitate liaisons among standards organizations.

Goals (1) and (2) in particular address aspects of ontology evaluation, while (3) and (4) will enable the adoption of the results of those goals.

4. NCOR: TOWARD ONTOLOGY EVALUATION

NCOR has begun to identify some of the questions we need to answer if we hope to enjoy well-defined ontology design techniques (quality of design), principled measurement methods (quality of evaluation), and ultimately higher quality ontologies (quality of content) in the future. Of course, in addition to the representational aspects of ontologies one must consider the role of automated reasoning technology: how is better reasoning over ontologies ensured, compared, promoted?

Some questions are:

- What are the appropriate languages for representing knowledge (ontologies, knowledge bases, rules, proofs, queries), and what are the desirable properties of these languages?
- What are the appropriate languages for run-time reasoning with knowledge?
- What should be the methodological principles and design criteria for developing the content of knowledge?
- What are the relative merits of competency questions, use cases, scenarios for guiding an ontology development effort? Can we use those competency questions and use cases for testing of the ontologies?
- How is knowledge content quality ensured? Should we certify the formal properties of content, e.g., soundness, completeness, consistency, accuracy, precision, and adherence to a principled methodology?
- How is knowledge content linked or aligned with other knowledge content, e.g., is a domain ontology situated in or linked to upper, middle, and reference ontologies?

The intent of such a program is to enable the mainstream adoption of ontologies, by describing and promoting principled approaches for developing and evaluating ontologies, and hence laying the foundation for a sound scientific and engineering discipline of ontologies. We intend to build upon the excellent work to date by groups such as the Knowledge Web Consortium [2] and as described in the survey by Brank et al. [3].

5. ONTOLOGY EVALUATION

Discussion of ontology evaluation methods within NCOR has focused so far on three potential non-exclusive thrusts, but all could be pursued in parallel: (1) development of an ontology and ontology tool competition or "bake-off," (2) principled certification of ontologies by a reviewing organization or community, and (3) the development of an ontology maturity model. The following sections describe these three thrusts.

5.1 An Ontology Evaluation Competition

This thrust involves development of an ontology evaluation competition along the lines of the US National Institute of Standards and Technology's (NIST) Text Retrieval Evaluation Conference (TREC) [4]. In the past, TREC has focused on different text-related evaluations: information retrieval, text summarization, and more recently digital video retrieval. The various evaluation tracks have included question-answering, cross-language retrieval, enterprise search, spam-filtering, and blogosphere behavior. In addition, NIST and Lockheed Martin have hosted the Information Interpretation and Integration Conference (I³CON) [5], which focused on the evaluation of ontology alignment tools, which was also the focus of the Evaluation of Ontology-based Tools Workshop (EON) in 2004 [6].

We envision this ontology evaluation competition as similar to the format of the (EON) Workshop [7], which focuses on the W3C Semantic Web languages of RDF(S), OWL, and the emerging rule languages. We think the ontology evaluation competition should expand to include ontologies represented in non-Semantic Web knowledge representation (KR) languages such as ISO Standard Common Logic (SCL) [8], the emerging Interoperable Knowledge Language (IKL) from the Interoperable Knowledge Representation for Intelligence Support (IKRIS) project [9], and even those in platform-specific KR languages such as Cyc's CycL [10] as long as those could be translated into one of the public standard languages.

The key questions to be answered in planning such a competition would include:

- What are some competitive tasks, challenges, or tracks that allow ontology content, ontology reasoning methods, and ontology tool development to better solve real-world problems and promote best practices?
- Where would the substantial resources for such an ontology evaluation competition come from?
- What would the organizational structure of such an evaluation competition look like, e.g., joint academic, commercial, governmental teaming? International standards consortium sponsorship?

- What metrics are important to consider when evaluating ontologies? Would such metrics clearly convey the value of a well-designed ontology?

5.2 Certification of Ontologies

An additional tack that could be pursued is to develop a principled certification process for ontologies. Such an approach might be considered a “Good Ontology Keeping” seal of approval for a certified ontology. An organization or community, centralized or distributed, membership-regulated or open, would be engaged to review submitted ontologies according to a prescribed set of evaluative properties and within a prescribed principled methodology. The result for a given ontology could be a grade or score for the quality of the ontology. Additional measures associated with an ontology could be domain, breadth of application or coverage within that domain, average taxonomic depth and relational density of nodes, completeness of axiomatic specification, etc. The methodology that the specific ontology development followed could also be described, e.g., Methontology [11].

5.3 Ontology Maturity Model

Another approach toward evaluating ontologies is to develop an ontology maturity model, perhaps along the lines of the Software Engineering Institute’s Capability Maturity Model Integration [12], i.e., gradations and decision procedures for maturity of ontologies, prospectively for both content and tools.

Such a model could establish levels of maturity that are defined by properties, such as degree of logical formalization, axiomatizability and satisfiability measures, strictness and properties of the ontology development process (competency questions, use cases, scenarios); degree and kinds of ontology documentation (metadata annotation, natural language descriptions of the concepts); references to authoritative subject matter sources; embedding or linking to or use of reference, utility, middle, and upper ontologies; application usage (areas); tool support (ontology development tools, runtime reasoning tools, support of deduction, induction, abduction, probabilistic reasoning, etc.).

6. CONCLUSIONS

This position paper describes some ideas and discussions that NCOR members have had recently concerning the prospects for ontology content and tool evaluation. It has proposed some questions that should be considered for ontology evaluation, and has suggested three prospective approaches for ontology evaluation: (1) the development of an ontology evaluation competition (2) certification of ontologies, and (3) development of an ontology maturity model.

We as NCOR members believe that ontology evaluation should be extended to include all ontologies that are represented in or translatable to a standard knowledge representation language, Semantic Web-based or not. Content, development methodology, automated reasoning aspects, and tool support are all important

for making the development and use of ontologies into a true scientific and engineering discipline.

7. ACKNOWLEDGMENTS

We note that the views expressed in this paper are those of the authors alone and do not reflect the official policy or position of The MITRE Corporation, Lockheed Martin Advanced Technology Laboratories, the National Institute of Standards and Technology, the National Center for Ontological Research, or any other company or individual.

This publication was prepared by United States Government employees as part of their official duties and is, therefore, a work of the U.S. Government and not subject to copyright.

8. REFERENCES

- [1] National Center for Ontological Engineering (NCOR) Inaugural Event. 2005. October 26, 2005. <http://ncor.buffalo.edu/inaugural/index.html>.
- [2] Hartmann, J. et al., Methods for ontology evaluation, http://www.aifb.uni-karlsruhe.de/Publikationen/showPublikation?publ_id=1021
- [3] Brank, J., Grobelnik, M. and Mladenic, D., A survey of ontology evaluation techniques, <http://kt.ijs.si/dunja/sikdd2005/Papers/BrankEvaluationSiKDD2005.pdf>
- [4] Text Retrieval Evaluation Conference (TREC). <http://trec.nist.gov/>.
- [5] Information Interpretation and Integration Conference (I³CON). <http://www.atl.external.lmco.com/projects/ontology/i3con.html>.
- [6] Evaluation of Ontology-based Tools (EON) 2004. <http://km.aifb.uni-karlsruhe.de/ws/eon2004>.
- [7] Evaluation of Ontologies for the Web (EON) 2006. <http://www.acm.org/sigs/pubs/proceed/template.html>.
- [8] ISO Standard Common Logic (SCL). <http://cl.tamu.edu/>.
- [9] Interoperable Knowledge Representation for Intelligence Support (IKRIS). <http://nrrc.mitre.org/NRRC/ikris.htm>.
- [10] CycL: the Cyc Representation language. http://www.cyc.com/technology/whatis_cyc_dir/howdoes_cyc_reason.
- [11] Fernandez M, Gomez-Perez A. et al. METHONTOLOGY: From ontological art towards ontological engineering. Proc AAAI-97 Spring Symposium Series on Ontological Engineering, Stanford, USA, pages 33-40.
- [12] Capability Maturity Model Integration, Software Engineering Institute, Carnegie-Mellon University. <http://www.sei.cmu.edu/cmmi/>.