Question Driven Semantics Interpretation for Collaborative Knowledge Engineering and Ontology Reuse

Khalid Latif, Edgar Weippl, A Min Tjoa
Institute of Software Technology & Interactive Systems
Vienna University of Technology,
Favoritenstrasse 9-11/188,
A-1040 Vienna, Austria
{klatif, weippl, amin}@ifs.tuwien.ac.at

Abstract

Ontology integration, alignment, and reuse are at the heart of Semantic Web vision. The alignment between two ontologies can be achieved easily provided both share the same axiomatic space. This in turn can be achieved by reusing constructs from foundational ontologies and standard domain thesauri. Although the current breed of ontology development methodologies and tools have made it a lot easier to build a new ontology, reusing concepts from existing ontologies remains difficult. The main reasons for this are the ambiguities in semantics interpretations of concepts from foundational ontologies, which are biased by philosophical aspirations, domain nuance, and design rationale introduced at the time of their creation. In this paper we have introduced a novel methodology for sense disambiguation to help the human expert in semantics interpretation. It is a collaborative and interaction intensive question driven approach based on a DOLCE aligned form of WordNet – OntoWordNet. The approach was used successfully for ontology reuse during the collaborative ontology building process; results are provided here.

1. Introduction

Recent efforts to realize a semantic web have accelerated research on the development of ontologies, a development that has been progressing slowly ever since the early efforts of the ancient Greek philosophers. Constrained by strong philosophical foundations and varying schools of thought, ontology development is still a difficult task for *non-ontologists*. One critical issue is to identify philosophical standing of the ontology. Making this determination at the outset of the development process can potentially constrain thinking, leading to an inadequate or incomplete definition, and ultimately may prove to be formally wrong [23] in satisfying a particular scenario. On the other, hand reusing an existing foundational ontology as a base could provide taxonomic and axiomatic contexts for the ontology. Ontologies are catalysts for knowledge sharing [3] and mediation in heterogeneous environments [17]. The use of foundational ontologies can further facilitate mutual understanding and interoperability among ontologies that vary otherwise.

In the ontology building process, identification of the concepts and patterns that should be modeled in the ontology is the most important and critical question. Adopting a high level view from upper ontologies provides an enormous jump start in answering this question [18]. Recently, different ontology development methodologies have emerged [14], some of which advocate the reuse of concepts or patterns from upper level (foundational) ontologies [2, 6]. The foundational ontologies are based on sound axiomatic theories and therefore have higher abstraction and formality level [20]. Scope of such ontologies may be specific to a domain, in which case they are characterized as core domain ontologies (see Figure 1).

Although the current breed of ontology management tools have made it a lot easier to build new ontologies, it is difficult to reuse concepts from existing ontologies. The main reason for this are ambiguities in semantics' interpretation of language [12] and concepts that are biased by philosophical orientation, domain nuance, and design constructs introduced at the time of their modeling. In line with the interaction paradigm – *knowledge should be confirmed by experience of actual perceptions that determine knowledge* [8] – ontologies should be built by human experts. Rather than automating the alignment task in ontology development, systems should be built for supporting human experts in alignment and reuse in the ontology building process [4].

1-4244-1500-4/07/$25.00 ©2007 IEEE

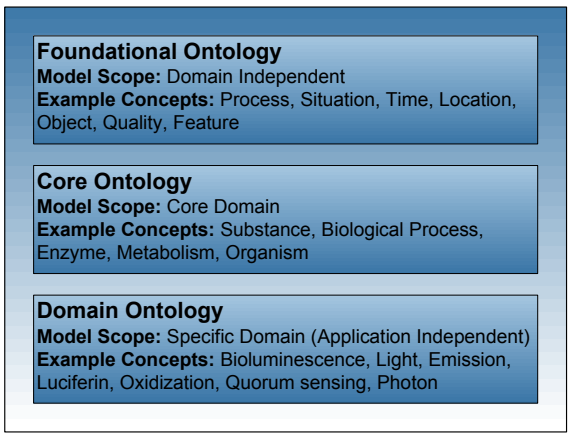

Figure 1. Classification of Ontologies.

Our hypothesis is that the concepts harvested during the ontology engineering workflow should be made available to the *Community of Practice* for rationalizing semantics based on the experience of actual perceptions that determine knowledge guided by the best practices followed in building foundational ontologies and thesauri. Effectively, top (level) concepts are aligned to concepts in foundational ontology, thus reusing its axiomatic context. In this paper we introduce a novel approach for realizing this hypothesis, which also provides necessary evidence to prove its expedience in ontology engineering. Our approach focuses on helping domain experts, rather than ontology experts, in eliciting domain knowledge by aligning the terminology with the foundational ontology. Helping domain experts through questions and the consequences of their answers, in finding right axiomatic context for the concepts, has resulted in effective semantics interpretation and ontology reuse in collaborative settings. After studying different foundational ontologies [1], we have chosen to use DOLCE in this work.

This paper is organized as follows: An outlook of existing methodologies emplying reuse of foundational ontologies is presented in Section 2. Within that section we also analyze different risk factors that might hinder the reuse of foundational ontologies. The interaction intensive, question driven approach for semantics interpretation and term alignment is elaborated in Section 3, followed by the details of collaboration support and workbench implementation in Section 4. Finally, we provide our findings and present an outlook for future work.

2. Foundational Ontologies and Reuse

In this section, we will present building blocks of ontology alignment, reuse and a brief overview of existing efforts related to collaborative ontology development involving reuse of foundational ontologies. The work presented here is very similar to ontology alignment, but an overview of existing efforts in the area is out of scope of this paper. For comprehensive details on matching, mapping, and alignment, refer to [15] and [21].

2.1. DOLCE & OntoWordNet

DOLCE is an ontology of *particulars* having four top level concepts: *endurant*, *perdurant*, *quality*, and *abstract*. It aims at "capturing the ontological categories underlying natural language and human commonsense." The work related to ontology design patterns [6] complements DOLCE ontology. CODeP are essentially *interconnected fragments* of this foundational ontology.

On the other hand, WordNet is a lexical database [5] and is used extensively by ontology authors to ground their ontologies [19]. To benefit from its coverage of terminology, it was aligned with DOLCE ontology [7]. The work resulted in a major restructuring of WordNet's top level concepts. For example, the WordNet synset $\langle process,\ physical\,process\rangle$ has hypernym (a.k.a. is a *kind of*) $\langle physical\,entity\rangle$, whereas OntoWordNet classifies it as $\langle phenomenon\rangle$. WordNet verb classes, compared to nouns in OntoWordNet, were investigated in [9] for semantics interpretation.

Reusing CODePs, even in the presence of OntoWordNet, is sometimes difficult because of ambiguities in the correct interpretation of domain terminology. Consider the following scenario: the concept *cooking* is modeled as a specialization of *edns:situation* in DOLCE and could easily be confused with *edns:activity*, or even *dol:process*, by the ontology author. Our work focuses on reusing CODePs by aligning domain terminology to correct DOLCE classes. The domain experts are guided through a question driven mechanism to disambiguate any confusion in interpretation of their terminology, in line with DOLCE.

2.2. Risk Factors

Indecision in interpretation of the semantics of concepts from foundational ontology may hinder the domain ontology development, especially in terms of the reuse of patterns and principles of the foundational ontology. We anticipate the following possible factors that may arise and hinder the domain ontology development process while using an upper level ontology, and give a brief account of their analysis against DOLCE and OntoWordNet.

Abstraction Level: *The upper ontologies do not reach down to the domain level*. Use of OntoWordNet can effectively resolve this issue. For example, the concept *photosynthesis* is described as *synthesis of compounds*

with the aid of radiant energy in WordNet. Its hypernym[1] chain leads to $\langle physical\ entity \rangle$, which adds to the confusion in finding correct axiomatic space for this concept in DOLCE. In contrast, OntoWordNet has redefined *photosynthesis* as a phenomenon, which is aligned with *edns:phenomenon* in DOLCE. Thus, the OntoWordNet mappings helped us find its correct alignment with DOLCE, which would otherwise be difficult due to abstraction in upper level foundational ontologies, in this case, DOLCE.

Formality Level: *We can't find a sufficient set of semantic descriptors to ontologize the domain terminologies in coherence with upper level ontology* - therefore failing to align the domain vocabulary with the upper level ontologies. The upper level ontology may have some formal philosophical assertions that cannot be matched with the specific domain ontology or which do not match the purpose and scope of the ontology. Although this issue can be resolved, to some extent, by using OntoWordNet vocabulary and its mapping with DOLCE, more examination of the issue and further research is needed. For example, modeling "chess game" using DOLCE vocabulary is difficult due to issues involving integration of Description Logics with Constraint Programming formalism.

Monolithic View: Although adopting a high level view from a single monolithic ontology is easier from the modeling point of view, it may still hinder the cause of interoperability. Major ontology players, such as John Sowa, advocate the use of multiple foundational ontologies coupled with mappings *to move along the lattice of different ontological commitments*[2]. The work done in [18] is a good reference point for easy and rigorous comparisons among different ontological approaches; concluding that the most important challenge for the Semantic Web is careful isolation of fundamental ontological commitments and their formal relationships. The use of lexical semantics could be investigated to further develop a system of semantic equivalence between two different foundational ontologies, for example, for DOLCE and SUMO.

Semantic Enrichment: The ontological enrichment of terminologies is possible, but the *semantic expressiveness of the resulting knowledge representation remains vacuous*. Prior research in [22] and [6] have proven that ontological enrichment of terminologies using well established principles from foundation ontologies not only contributes to semantic expressiveness but also helps in achieving on-the-fly ontology matching.

[1] Hypernym is the more general class of another synset.

[2] http://suo.ieee.org/email/msg03804.html

2.3 Collaborative Ontology Matching

Collaborative ontology development and mapping has been investigated in different research projects such as in [13] and [24]. The work presented in [24] is based on the assumption that *alignments established by an individual are most likely valid and valuable*. This hypothesis is contrary to our approach in that the ontology authors are guided in establishing valid alignments by either reusing fragments of foundational ontologies or grounding their vocabulary in the foundational ontology and following already established modeling constructs. An interactive question driven approach can help the individual in correct interpretation of the concepts for achieving valid alignments and correct reuse.

3. Question Driven Semantics Interpretation

The workflow of our semantics interpretation methodology starts with identifying the purpose and eliciting a description and usage of the intended ontology. The ontology usage scenarios are structured and include different sections, such as *problem statement*, *purpose*, and *process description*. Describing the purpose and problem statement at early stages of ontology building is in line with different ontology development methodologies [11, 14]. Subsequently, term extraction is applied to the domain corpus including documents and existing terminology databases. Selected terms are further analyzed and domain experts are guided in aligning these terms with DOLCE. Complete details of the DynamOnt methodology are present in [2, 10]. Here we will assume a *cooking information system* scenario and will elaborate the details of question driven methodology for terminology alignment.

3.1. Bottom-up Analysis

The alignment procedure starts with bottom-up analysis. For each domain concept C, its possible mappings $\mathcal{S}_c = \{Syn_1, Syn_2, \ldots Syn_n\}$ with WordNet synsets are discovered and presented to the domain expert.

Now, alignment of each synset $Syn_i \in \mathcal{S}_c$ with corresponding DOLCE class (a *subclass* of *dol:particular*) is identified by using OntoWordNet and each is represented as another set $\mathcal{P}_{syn} = \{dolp_1, dolp_2, \ldots dolp_n\}$ where $f : (Syn_i \in \mathcal{S}_c) \rightarrow (dolp_i \in \mathcal{P}_{syn})$. The set $\mathcal{P}_{syn}$ renders different possibilities for aligning the concept C with DOLCE's axiomatic space. Two special situations may arise and need to be processed. In the first case, elements of $\mathcal{P}_{syn}$ may be equal, which means that all WordNet synsets in $\mathcal{S}_c$ were aligned to the same DOLCE class. In the second case, $\mathcal{P}_{syn}$ might have only one element. This situation arises when the domain expert selects only single WordNet

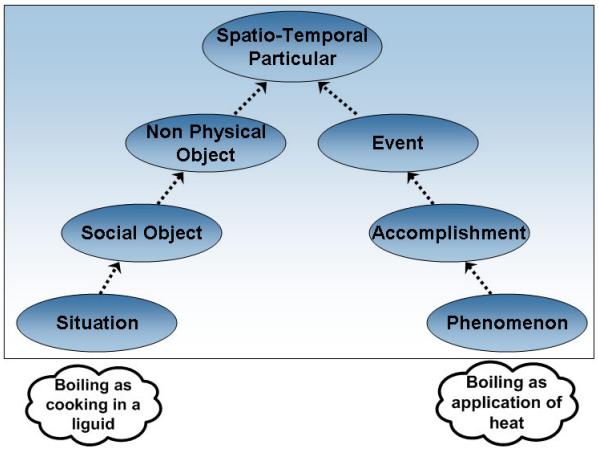

Figure 2. Fragment of DOLCE taxonomy and alignment for different senses of *boiling*.

sense as being relevant for concept C. In both cases, the concept C is aligned with the first element in $\mathcal{P}_{syn}$ without proceeding further. Otherwise, we proceed with the normal flow of the alignment methodology.

For example, consider the concept *boiling*. The noun $\langle boiling \rangle$ has two senses 1) *the application of heat* and 2) *cooking in a liquid*. These senses are respectively interpreted as $\langle natural\,process \rangle$ and $\langle change\,of\,state \rangle$ by WordNet and aligned with *edns:phenomenon* and *edns:situation* in DOLCE. Consequently, DOLCE's class hierarchy is traversed to determine decision points – the places of deviations in the synset alignment with DOLCE for varying senses (see Figure 2).

3.2. Class Hierarchy as Concept Chains

Each path of the class hierarchy from the previous step is transformed to a concept chain. Concept chains are needed for efficient comparisons and traversing the class hierarchy. A concept chain is an ordered set of concepts based on their *sub-class* relationship. Concept chains are virtual collections and support navigation through operations, such as *next*, *previous*, which are delegated to the actual taxonomy. The first element in the set is the leaf class and the last element is the root class in the class hierarchy. In general, for an ontology $\mathcal{O}$ a concept chain Φ_{c1} for a concept C_1 is defined as follows:

$$\Phi_{c1} = \langle C_1, C_2, \cdots C_r \rangle \quad (1)$$

such that the concept $C_i, C_{i+1} \in \Phi_{c1}$ satisfies (`subclass-of` C_i C_{i+1}). From the previous boiling example we realize two concept chains for the two senses.

Φ_{boil1} = ⟨ *phenomenon, accomplishment, event, perdurant, spatio-temporal-particular, particular*⟩

Φ_{boil2} = ⟨ *situation, non-agentive-social-object, social-object, non-physical-object, non-physical-endurant, endurant, spatio-temporal-particular, particular* ⟩

From the concept chains, it is evident that both alignments contradict at the point of further classification of *dol:spatio-temporal-particular*. Their classification as being an *dol:event* or *dol:non-physical-object*[3] (see Figure 2) is sorted out and domain experts are aided in deciding on the correct alignment for the target domain ontology. In the next section we will explain our question driven approach for resolving such ambiguities in alignment.

3.3. Question and Answer Model

Question answering using Semantic Web technologies is not a new idea. PowerAqua is one such system designed to make use of distributed semantic contents to answer user queries in natural language [16]. Our approach begins the other way around entirely – the DynamOnt system asks questions of the domain expert rather than the user asking questions. Although the questions are posed in natural language, consequences of their answers are first semantically described.

Questions correspond to concepts in the foundational ontology, in our case, DOLCE. The structure of the question model is explained below:

q-for Each question corresponds to a specific DOLCE class referred to by this slot.

description This slot describes the body text of the question and is taken, for the most part, from DOLCE's description of the class. Some modifications are made in order to make it easily understandable for the domain experts. This description also includes a variable $concept$ which is replaced with the user term.

hint It provides an exemplar to help the domain expert in answering the question.

It is worth mentioning that harvesting competency questions from DOLCE turned out to be the most difficult task in the implementation of our methodology. Initially we only modeled 25 questions against different DOLCE classes (for the most part, those mentioned in CODePs [6]).

Five possible answers are permitted for each question to declare the consent for aligning the user term with the

[3] We observed that most of the contradictions involve *edns:situation* and `edns:phenomenon`. For this reason, we have deliberately presented a similar case, although there are other intriguing cases of alignment.

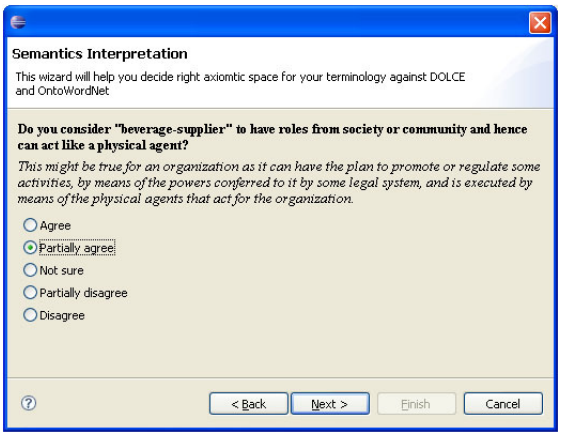

Figure 3. A screenshot from terminology alignment wizard

DOLCE concept referred to by the question. The answers include 1) *Agree* – relevant DOLCE concept, 2) *Partially agree* – agreement with some uncertainty, 3) *Not sure* – user is not sure about the answer, 4) *Partially disagree* – disagreement with some uncertainty, and 5) *Disagree* – the DOLCE concept is not a right match for the user term. Options 2 and 4 are included to incorporate weaker notion of (dis-)agreement. Each answer is weighted symmetrically, that is, an agreement or disagreement gets equal weight.In addition, relative weights are allowed by introducing variables α as a factor of agreement and β as a factor of uncertainty. Initially they are set to 3 and 0.4 respectively.

3.4. Top-down Analysis

For each concept chain Φ_c, a corresponding answer set $\Phi_A = \langle A_1, A_2, ...A_r \rangle$ is constructed such that $A_i \in \Phi_A$ is a relevance score for class $C_i \in \Phi_c$. The relevance score is computed from users' answers for the questions against the DOLCE alignment. Initially, elements of the answer set are initialized with zeros.

The concept chains correspond to hierarchical paths in DOLCE. The paths are established after bottom-up analysis of WordNet synsets to DOLCE alignment. In the next phase the concept chains are traversed in reverse order. It is a top-down approach considering the hierarchy of the classes in a concept chain . The domain experts are asked questions about each class excluding those questions that don't necessarily add new knowledge. For example in *boiling* example, asking the domain expert if boiling could be classified as *dol:spatio-temporal-particular* doesn't resolve any ambiguity. This strategy reduced the number of questions required to effectively align the term with a DOLCE class. A question against the leaf class is also asked to confirm the alignment (c.f. Figure 3). Finally, the user's answer scores for all concept chains are enumerated and the concept chain with the highest score wins the alignment decision. The algorithm for processing two concept chains Δ_{c1} and Δ_{c2} is presented below:

Algorithm 1 ***Processing Two Concept Chains for Questions***
This algorithm traverses two concept chains Φ_{c1} and Φ_{c2}, having corresponding answer sets Ψ_{A1} and Ψ_{A2}, and asks competency questions of the domain expert for aligning the concept C with DOLCE.

1. / set the pointer to last (root) item in the chain */*
2. $\Phi_{c1}.MoveLast()$
3. $\Phi_{c2}.MoveLast()$
4. / Skip till contradiction */*
5. ***while*** $\Phi_{c1}.Current() = \Phi_{c2}.Current()$ ***do***
6. $\Phi_{c1}.Move(-1)$
7. $\Phi_{c2}.Move(-1)$
8. ***end while***
9. ***repeat***
10. / Compute scores and get the concept chain and corresponding answer set with relatively higher score this far*/*
11. $[\Phi_{cx}, \Psi_{Ax}] \leftarrow$ *ComputeScore(*$[\Phi_{c1}, \Psi_{A1}], [\Phi_{c2}, \Psi_{A2}]$*)*
12. / Get concept at the current index */*
13. $[cindex, C_x] \leftarrow \Phi_{cx}.Current()$
14. ***if*** $cindex = -1$ ***then***
15. / We have reached the end of the concept chain but have only achieved partial agreement or no agreement at all. Ask the user to either align with the class corresponding to the concept chain having a relatively higher score or restart the procedure after selecting different senses of C from WordNet.*/*
16. ***break***
17. ***end if***
18. $score \leftarrow AskQuestionFor(C_x, C)$
19. $\Psi_{Ax}[cindex] \leftarrow score$
20. $\Phi_{cx}.Move(-1)$
21. ***until*** $\sum \Psi_{Ax} < \tau$ */* τ *is agreement threshold*/*

3.5. Variation in Top-down Analysis

Some concepts are more easily aligned with DOLCE than others. *Cooking*, for instance, has only one sense as a noun in WordNet and is aligned with *edns:situation*. The only competency question required was for the leaf class, *edns:situation* in this case, and the alignment was achieved successfully.

An interesting scenario is when there is only a single possible alignment with DOLCE but the domain expert states otherwise. For example, OntoWordNet aligns the

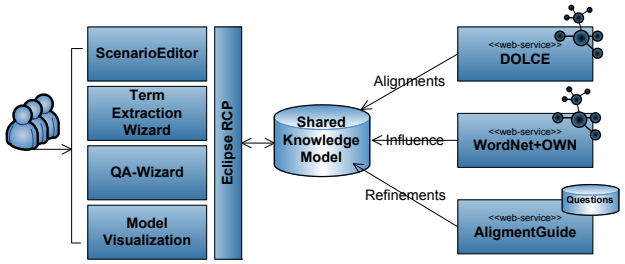

Figure 4. Workbench Architecture using a combination of Eclipse RCP and Web Services

concept *recipe* with DOLCE as *edns:situation*. Its concept chain is as follows: ⟨*situation non-agentive-social-object, social-object, non-physical-object, non-physical-endurant, endurant, spatio-temporal-particular, particular*⟩. The domain expert is asked to confirm the alignment with *edns:situation*. As a consequence of disagreement, a question is asked for each class in a bottom-up way to find out the top most class with which user agrees to align. Alternate classifications of that class are then traversed to find the right match. In the case of *recipe*, it turned out that the user was more interested in *recipe* being a *edns:plan* identified after being asked the question about sub classes of *edns:non-agentive-social-object*.

4. Implementation and Collaboration Support

To support collaborative ontology building and terminology alignment, the backend of the workbench consists of different web-services. The client interface, on the other hand, was developed as an Eclipse Rich Client Platform[4]. WordNet database and DOLCE & OntoWordNet ontologies are provided as web-services[5]. With current Semantic Web frameworks and APIs, such as Jena, it is hard to load both OntoWordNet and DOLCE on one machine along with the DynamOnt workbench because of mammoth memory requirements. For such pragmatic reasons we have deployed these web-services, in a distributed setting, on separate machines. To improve efficiency, the OntoWordNet web-service uses LWWN[6], which is a Lucene[7] index of the WordNet database.

An overview of the workbench architecture could be seen in Figure 4 and details of operations supported by the OntoWordNet web-service are given below.

GetWordSenses: Given a term returns URI's of all the matching WordNet.

GetDescription: Given a WordNet term this operation returns its detailed description. This operation also uses LNWN for its internal operations.

GetParent: Returns immediate parent of the WordNet term from OntoWordNet.

GetSynsetChain: Returns hierarchical chain of parents (subClass relations).

GetHypernymChain: Returns hierarchical chain of hypernyms of a particular WordNet synset.

GetAlignment: Returns DOLCE alignment of a WordNet synset from the OntoWordNet mappings.

Each user is allowed to select a concept and align it with the relevant DOLCE class by answering the questions posed by the system. Answer sets for each user are accumulated to match corresponding concept chains. For example, for three concept chains and 10 users participating in the alignment, there will be 3 distinct collections each containing 10 answer sets. The answer set collection with the highest accumulated score is adopted for the alignment decision and the concept is aligned with the leaf concept referred to in the concept chain for that answer set collection.

5. Conclusion and Future Work

Automatic semantic matchmaking is a challenge for the Semantic Web in general and for realizing automatic discovery of Semantic Web Services on a large scale. Currently, ontologies are developed mainly to implement software systems that focus on specific problems, without considering the ontology reuse and alignment aspects. The focal point of such ontologies is their usability and not the soundness of axiomatic theories.

The trade-off between usability and formality is a difficult one. On the one hand, formality comes with increased complexity, making it hard for current inference tools to interpret the semantics. On the other hand, lightweight taxonomic ontologies grounded on best practices and developed by reusing fragments of foundational ontologies can achieve formality without compromising usability.

In this paper we have presented a methodology for building formal ontologies by aligning the domain terminology with the foundational ontology. Reusing the axiomatic context of the foundational ontology effectively resolves misconstructions in the domain modeling process. Competency questions guide the domain experts in terminology alignments.

[4] http://www.eclipse.org/rcp/
[5] http://storm.ifs.tuwien.ac.at:8081/
[6] http://eden.dei.uc.pt/ nseco/lwwn.tar.gz (visited September 2006)
[7] http://lucene.apache.org

We have planned to evaluate our methodology from the perspective of users; for example, how easy it is for a domain expert to understand the generated question. And further, does OntoWordNet and DOLCE cover the needs of several domains, and are there domains which do not benefit from this methodology.

6. *

Acknowledgments

The authors would like to acknowledge the support of the DynamOnt team. The work presented in this paper is result of various group discussions and is funded by the Austrian Federal Ministry for Transport, Innovation, and Technology under the FIT-IT contract FFG 809256/5512 and also by industrial partners factline Webservices and ProCom Strasser. We would also like to thank Aldo Gangemi for the extended help and support he gave to the DynamOnt project members on DOLCE, OntoWordNet, and Ontology Design Patterns.

References

[1] Upper ontologies with specific consideration of DOLCE, SUMO and Sowa's upper level ontology. Deliverable 121, DynamOnt Project, August 2005. http://dynamont.factlink.net.

[2] W. Behrendt, E. Gahleitner, K. Latif, A. Gruber, and R. Westenthaler. Specification of methodology and workbench for dynamic ontology creation. Deliverable 201, DynamOnt Project, 2006. http://dynamont.factlink.net.

[3] T. Edgington, B. Choi, K. Henson, T. Raghu, and A. Vinze. Adopting ontology to facilitate knowledge sharing. *Communications of ACM*, 47(11):85–90, 2004.

[4] S. M. Falconer, N. F. Noy, and M.-A. Storey. Towards understanding the needs of cognitive support for ontology mapping. In *Proceedings of International Workshop on Ontology Matching (OM-2006)*, Georgia, USA, November 2006.

[5] C. Fellbaum, editor. *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge, MA, 1998.

[6] A. Gangemi. Ontology design patterns for semantic web content. In *Proceedings of 4th International Semantic Web Conference (ISWC)*, volume 3729 of *LNCS*, pages 262–276, Galway, Ireland, November 2005. Springer.

[7] A. Gangemi, R. Navigli, and P. Velardi. The OntoWordNet project: extension and axiomatization of conceptual relations in WordNet. In *On The Move to Meaningful Internet Systems*, volume 2888 of *LNCS*, pages 820–838, Catania, Italy, November 2003. Springer.

[8] D. Goldin and P. Wegner. Principles of interactive computation. In D. Goldin, S. A. Smolka, and P. Wegner, editors, *Interactive Computation - The New Paradigm*, pages 25–37. Springer, 2006.

[9] F. Gomez. An algorithm for aspects of semantic interpretation using an enhanced WordNet. In *Proceedings of the 2nd Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL)*, CMU, Pittsburgh, June 2001.

[10] A. Gruber, R. Westenthaler, and E. Gahleitner. Supporting domain experts in creating formal knowledge models. In *Proceedings of I-KNOW'06*, Graz, Austria, September 2006.

[11] T. Gruber. Toward principles for the design of ontologies used for knowledge sharing. *Human-Computer Studies*, 43(1):907–928, 1995.

[12] D. Harel and B. Rumpe. Meaningful modeling: What's the semantics of "semantics?". *IEEE Computer*, 37(10):64–72, October 2004.

[13] C. Holsapple and K. D. Joshi. A collaborative approach to ontology design. In *Communications of ACM*, volume 45(2), pages 42–47, 2002.

[14] D. Jones, T. Bench-Capon, and P. Visser. Methodologies for ontology development. In *Proceedings of IFIP XV IT & KNOWS*, pages 62–75, Budapest, Hungary, 1998.

[15] Y. Kalfoglou and M. Schorlemmer. Ontology mapping: The state of the art. *The Knowledge Engineering Review*, 18, 2003.

[16] V. Lopez, E. Motta, and V. Uren. Poweraqua: Fishing the semantic web. In *Proceedings of the 3rd European Semantic Web Conference (ESWC)*, volume 4011 of *LNCS*, pages 393–410, Budva, Montenegro, June 2006. Springer.

[17] O. Lyttleton, D. Sinclair, and D. Tracey. Mediating between heterogeneous ontologies using schema matching techniques. In *IEEE International Conference on Information Reuse and Integration (IRI)*, pages 247–252, Las Vegas, Nevada, August 2005. IEEE SMC Society.

[18] C. Masolo, S. Borgo, A. Gangemi, N. Guarino, and A. Oltramari. The WonderWeb library of foundational ontologies. Deliverable 18, WonderWeb Project, 2003.

[19] I. Niles and A. Pease. Linking lexicons and ontologies: Mapping WordNet to the Suggested Upper Merged Ontology. In *Proceedings of International Conference on Information and Knowledge Engineering*, Las Vegas, Nevada, 2003. IEEE.

[20] S. Schaffert, A. Gruber, and R. Westenthaler. A semantic wiki for collaborative knowledge formation. In *Proceedings of Semantics*, Vienna, Austria, November 2005.

[21] P. Shvaiko and J. Euzenat. A survey of schema-based matching approaches. *Journal on Data Semantics*, IV:146–171, 2005.

[22] B. Smith and C. Rosse. The role of foundational relations in the alignment of biomedical ontologies. In *Proceedings of 11th World Congress on Medical Informatics (MedInfo)*, volume 107 of *Studies in Health Technology and Informatics*, pages 444–448, San Francisco, CA, August 2004. IOS Press.

[23] M. Uschold. Building ontologies: Towards a unified methodology. In *Proceedings of 16th Annual Conference of the British Computer Society Specialist Group on Expert Systems*, Cambridge, UK, December 1996.

[24] A. Zhdanova and P. Shvaiko. Community-driven ontology matching. In *Proceedings of the 3rd European Semantic Web Conference (ESWC)*, volume 4011 of *LNCS*, pages 34–49, Budva, Montenegro, June 2006. Springer.