

Semantic Web: The Story So Far

Ian Horrocks
University of Manchester
Manchester, UK
horrocks@cs.man.ac.uk

ABSTRACT

The goal of Semantic Web research is to transform the Web from a linked document repository into a distributed knowledge base and application platform, thus allowing the vast range of available information and services to be more effectively exploited. As a first step in this transformation, languages such as OWL have been developed; these languages are designed to capture the knowledge that will enable applications to better understand Web accessible resources, and to use them more intelligently. Although fully realising the Semantic Web still seems some way off, OWL has already been very successful, and has rapidly become a de facto standard for ontology development in fields as diverse as geography, geology, astronomy, agriculture, defence and the life sciences. An important factor in this success has been the availability of sophisticated tools with built in reasoning support. The use of OWL in large scale applications has brought with it new challenges, both with respect to expressive power and scalability, but recent research has also shown how the OWL language and OWL tools can be extended and adapted to meet these challenges.

1. INTRODUCTION

While phenomenally successful in terms of size and number of users, today's World Wide Web is fundamentally a relatively simple artefact. Web content consists mainly of distributed hypertext, and is accessed via a combination of keyword based search and link navigation. This simplicity has been one of the great strengths of the Web, and has been an important factor in its popularity and growth: naive users are able to use it, and can even create their own content.

The explosion in both the range and quantity of Web content has, however, highlighted some serious shortcomings in the hypertext paradigm. In the first place, the required content becomes increasingly difficult to locate using the search and browse paradigm. Finding information about people with very common names (or with famous namesakes) can, for example, be a frustrating experience. More

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

W4A2007 Keynote, May 07–08, 2007, Banff, Canada. Co-Located with the 16th International World Wide Web Conference.

Copyright 2007 ACM 1-59593-590-8/06/0010 ...\$5.00.

complex queries can be even more problematical: a query for “animals that use sonar but are neither bats nor dolphins”¹ may either return many irrelevant results related to bats and dolphins (because the search engine failed to understand the negation), or may fail to return many relevant results (because most relevant Web pages *also* mention bats and/or dolphins). More complex tasks may be extremely difficult, or even impossible. Examples of such tasks include locating information in data repositories that are not directly accessible to search engines [34], or finding and using so-called web services [24].

If human users have difficulty accessing web content, the problem is even more severe for automated processes. This is because web content is primarily intended for presentation to and consumption by human users: HTML markup is mainly concerned with layout, size, colour and other presentational issues. Moreover, web pages increasingly use images, often including active links, to present information. Human users are able to interpret the significance of such features, and thus understand the information being presented, but this may not be so easy for an automated process or “software agent”. It is easy to imagine that similar difficulties might be experienced by users with cognitive or sensory impairments.

The Semantic Web aims to overcome some of the above mentioned problems by making web content more accessible to automated processes; the ultimate goal is to transform the existing web into “... a set of connected applications ... forming a consistent logical web of data ...” [3]. This is to be achieved by adding *semantic annotations* to Web content, i.e., annotations that describe the meaning of the content.

In the following we will examine in a little more detail what semantic annotations will look like, how they describe meaning, and how automated processes can exploit such descriptions.

2. BACKGROUND

As we mentioned above, the key idea behind the semantic web is to explicate the *meaning* of web content by adding semantic annotations. If we assume for the sake of simplicity that such annotations take the form of XML style tags, we could imagine a fragment of a web page being annotated as follows:

¹Thanks to Uli Sattler for this example.

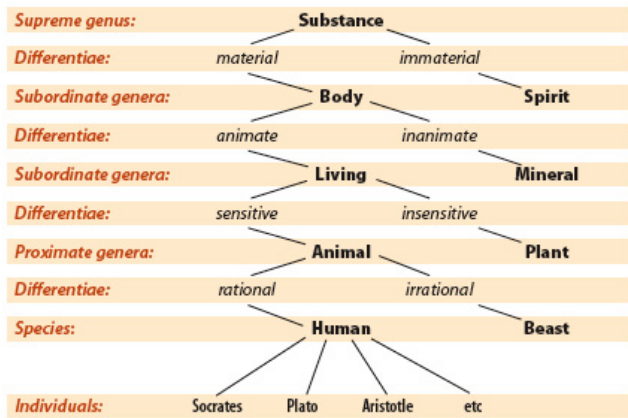


Figure 1: Tree of Porphyry.

`<Wizard>Harry Potter</Wizard>` has a pet called `<SnowyOwl>Hedwig</SnowyOwl>`.

Taken in isolation, however, such annotations are of only limited value: the problem of understanding the terms used in the text has simply been transformed into the problem of understanding the terms used in the labels. A query for information about raptors, for example, may not retrieve this text, even though owls are raptors. This is where *ontologies* come into play: they provide a mechanism for introducing a vocabulary and giving precise meanings to the terms in the vocabulary. A suitable ontology might, for example, introduce the term `SnowyOwl`, and include the information that `SnowyOwls` are kinds `Owl`, and that `Owls` are kinds of `Raptor`. Moreover, if this information is represented in a way that is accessible to our query engine, then it would be able to recognise that the above text is relevant to our query about raptors.

Ontology, in its original philosophical sense, is a fundamental branch of metaphysics focussing on the study of existence; its objective is to determine what entities and types of entities actually exist, and thus to study the structure of the world. The study of ontology can be traced back to the work of Plato and Aristotle, and from the very beginning included the development of hierarchical categorisations of different kinds of entity and the features that distinguish them: the well known “tree of Porphyry”, for example, identifies animals and plants as sub-categories of living things distinguished by animals being *sensitive*, and plants being *insensitive* (see Figure 1).

In computer science, an ontology is usually taken to be a model of (some aspect of) the world; it introduces vocabulary describing various aspects of the domain being modelled, and provides an explicit specification of the intended meaning of the vocabulary. This specification often includes classification based information not unlike that in Porphyry’s famous tree. For example, Figure 2 shows a screenshot of a `Pizza` ontology as displayed by the Protégé ontology design tool [22]. The ontology introduces various pizza related vocabulary (some of which can be seen in the left hand panel), such as “`NamedPizza`” and “`RealItalianPizza`”, and arranges it hierarchically: `RealItalianPizza` is, for example, a sub-category of `NamedPizza`. The other panels display information about the currently selected vocabulary term, `RealItalianPizza` in this case, describing its mean-

ing: a `RealItalianPizza` is a `Pizza` whose country of origin is Italy; moreover, a `RealItalianPizza` always has a `ThinAndCrispyBase`. Ontologies can be used to annotate and to organise data from the domain: if our data includes instances of `RealItalianPizza`, then we can return them in response to a query for instances of `NamedPizza`.

3. THE WEB ONTOLOGY LANGUAGE OWL

The architecture of the Web depends on agreed standards such as HTTP that allow information to be shared and exchanged. A standard ontology language is, therefore, a prerequisite if ontologies are to be used in order to share and exchange *meaning*. Recognising this fact, the World Wide Web Consortium (W3C) set up a standardisation working group to develop such a language. The result of this activity was the OWL ontology language standard [25]. OWL exploited existing work on languages such as OIL [8] and DAML+OIL [14] and, like them, was based on a Description Logic (DL). In the following we will briefly introduce DLs and OWL. For more complete information the reader should consult The Description Logic Handbook [1], and the OWL specification [25].

3.1 Description Logic

Description logics (DLs) are a family of logic-based knowledge representation formalisms; they are descendants of Semantic Networks [35] and KL-ONE [4]. These formalisms all adopt an object-oriented model, similar to the one used by Plato and Aristotle, in which the domain is described in terms of individuals, *concepts* (usually called *classes* in ontology languages), and *roles* (usually called relationships or properties in ontology languages). Individuals, e.g., “Socrates” are the basic elements of the domain; concepts, e.g., “Human”, describe sets of individuals having similar characteristics; and roles, e.g., “hasPupil” describe relationships between pairs of individuals, such as “Socrates hasPupil Plato”.

As well as *atomic* concept names such as `Human`, DLs also allow for concept descriptions to be composed from atomic concepts and roles. Moreover, it is possible to assert that one concept (or concept description) is subsumed by (is a sub-concept of), or is exactly equivalent to, another. This allows for easy extension of the vocabulary by introducing new names as abbreviations for descriptions. For example, using standard DL notation, we might write:

$$\text{HappyParent} \equiv \text{Parent} \sqcap \forall \text{hasChild} . (\text{Intelligent} \sqcup \text{Athletic})$$

This introduces the concept name `HappyParent`, and asserts that its instances are just those individuals that are instances of `Parent`, and all of whose children are instances of either `intelligent` or `athletic`.

Another distinguishing feature of DLs is that they are logics, and so have a formal semantics. DLs can, in fact, be seen as decidable subsets of first-order predicate logic, with individuals being equivalent to constants, concepts to unary predicates and roles to binary predicates. As well as giving a precise and unambiguous meaning to descriptions of the domain, this also allows for the development of reasoning algorithms that can be used to answer complex questions about the domain. An important aspect of DL research has been the design of such algorithms, and their implementation in (highly optimised) reasoning systems that can be

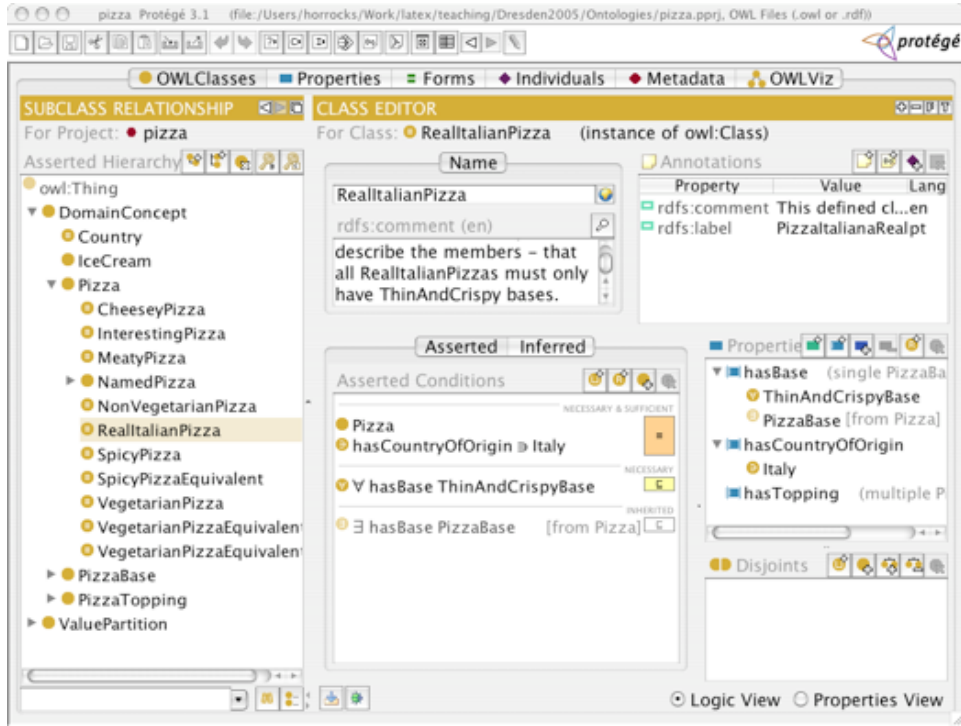


Figure 2: Example pizza ontology.

used by applications to help them “understand” the knowledge captured in a DL based ontology. We will return to this point in Section 4.

A given DL is characterised by the set of constructors provided for building concept descriptions. These typically include at least intersection (\sqcap), union (\sqcup) and complement (\neg), as well as restricted forms of existential (\exists) and universal (\forall) quantification, which in OWL are called, respectively, *someValuesFrom* and *allValuesFrom* restrictions. OWL is based on a very expressive DL called *SHOIN* that also provides cardinality restrictions (\geq , \leq) and enumerated classes (called *oneOf* in OWL) [15, 17]. Cardinality restrictions allow, e.g., for the description of a concept such as people who have at least two children, while enumerated classes allow for classes to be described by simply enumerating their instance, e.g.,:

$$\text{EUCountries} \equiv \{\text{Austria}, \dots, \text{UK}\}$$

SHOIN also provides for transitive roles, allowing us to state, e.g., that if x has an ancestor y and y has an ancestor z , then z is also an ancestor of x , and for inverse roles, allowing us to state, e.g., that if z is an ancestor of x , then x is also a descendent of z . The constructors provided by OWL, and the equivalent DL syntax, are summarised in Figure 3.

In DLs it is usual to separate the set of statements that establish the vocabulary to be used in describing the domain (what we might think of as the schema) from the set of statements that describe some particular situation that instantiates the schema (what we might think of as data); the former is called the TBox (Terminology Box), and the latter the ABox (Assertion Box). An OWL ontology is simply equivalent to a set of *SHOIN* TBox and ABox state-

Constructor	DL Syntax	Example
intersectionOf	$C_1 \sqcap \dots \sqcap C_n$	Human \sqcap Male
unionOf	$C_1 \sqcup \dots \sqcup C_n$	Doctor \sqcup Lawyer
complementOf	$\neg C$	\neg Male
oneOf	$\{x_1 \dots x_n\}$	{john, mary}
allValuesFrom	$\forall P.C$	\forall hasChild.Doctor
someValuesFrom	$\exists r.C$	\exists hasChild.Lawyer
hasValue	$\exists r.\{x\}$	\exists citizenOf.{USA}
minCardinality	$(\geq n r)$	$(\geq 2 \text{ hasChild})$
maxCardinality	$(\leq n r)$	$(\leq 1 \text{ hasChild})$
inverseOf	r^-	hasChild $^-$

Figure 3: OWL constructors

ments. This mixing of schema and data is quite unusual (in fact ontologies are usually thought of as consisting only of the schema part), but does not affect the meaning—from a logical perspective, *SHOIN* KBs and OWL ontologies are just sets of axioms.

The main difference between OWL and *SHOIN* is that OWL ontologies use an RDF based syntax intended to facilitate their use in the context of the Semantic Web. This syntax is rather verbose, and not well suited for presentation to human beings. E.g., the description of *HappyParent* given above would be written in OWL’s RDF syntax as follows:

```
<owl:Class>
  <owl:intersectionOf rdf:parseType="collection">
    <owl:Class rdf:about="#Parent"/>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#hasChild"/>
      <owl:allValuesFrom>
```

```

<owl:unionOf rdf:parseType=" collection">
  <owl:Class rdf:about="#Intelligent"/>
  <owl:Class rdf:about="#Athletic"/>
</owl:unionOf>
</owl:allValuesFrom>
</owl:Restriction>
</owl:intersectionOf>
</owl:Class>

```

4. ONTOLOGY REASONING

We mentioned in Section 3.1 that the design and implementation of reasoning systems is an important aspect of DL research. The availability of such reasoning systems was one of the motivations for basing OWL on a DL. This is because reasoning is essential in supporting both the design of high quality ontologies, and the deployment of ontologies in applications.

4.1 Reasoning at design time

Ontologies may be very large and complex: the well known Snomed clinical terms ontology includes, for example, more than 200,000 class names [31]. Building and maintaining such ontologies is very costly and time consuming, and providing tools and services to support this “ontology engineering” process is of crucial importance to both the cost and the quality of the resulting ontology. State of the art ontology development tools, such as SWOOP [19], Protégé [22], and TopBraid Composer (see <http://www.topbraidcomposer.com/>), therefore use a DL reasoner, such as FaCT++ [33], Racer [11] or Pellet [29], to provide feedback to the user about the logical implications of their design. This typically includes (at least) warnings about inconsistencies and redundancies.

An inconsistent (sometimes called unsatisfiable) class is one whose description is “over-constrained”, with the result that it can never have any instances. This is typically an unintended feature of the design—why introduce a name for a class that can never have any instances—and may be due to subtle interactions between descriptions. The ability to detect such classes and bring them to the attention of the ontology engineer is, therefore, a very useful feature.

It is also possible that the descriptions in the ontology mean that two classes necessarily have exactly the same set of instances, i.e., that they are alternative names for the same class. This may be desirable in some situations, e.g., to capture the fact that “Myocardial infarction” and “Heart attack” mean the same thing. It could, however, also be the inadvertent result of interactions between descriptions, and so it is also useful to be able to alert users to the presence of such “synonyms”.

In addition to checking for inconsistencies and synonyms, ontology development tools usually also check for implicit subsumption relationships, and update the class hierarchy accordingly. This is also a very useful design aid: it allows the ontology developer to focus on class descriptions, leaving the computation of the class hierarchy to the reasoner, and it can also be used by the developer to check if the hierarchy induced by the class descriptions is consistent with their intuition.

Recent work has also shown how reasoning can be used to support modular design [6] and module extraction [5], important techniques for working with large ontologies. When developing a large ontology such as SNOMED, it is useful

if not essential to divide the ontology into modules, e.g., to facilitate parallel work by a team of ontology developers. Reasoning techniques can be used to alert the developers to unanticipated and/or undesirable interactions between the various modules. Similarly, it may be desirable to extract from a large ontology a smaller module containing all the information relevant to some subset of the domain, e.g., heart disease—the resulting small(er) ontology will be easier for humans to understand and easier for applications to use. Reasoning can be used to compute a module that is as small as possible while still containing all the necessary information.

Finally, in order to maximise the benefit of all these services, a modern system should also be able to explain its inferences: without this facility, users may find it difficult to repair errors in the ontology and may even start to doubt the correctness of the reasoning system. Explanation typically involves computing a (hopefully small) subset of the ontology that still entails the inference in question, and if necessary presenting the user with a chain of reasoning steps [20].

4.2 Reasoning in deployment

Reasoning is also important when ontologies are deployed in applications—it is needed, e.g., in order to answer structural queries about the domain and to retrieve data. If we assume, for example, an ontology that includes the description of HappyParent given in Section 3.1, and we know that John is a HappyParent, that John has a child Mary (i.e., John hasChild Mary), and that Mary is not Athletic, then we would like to be able to infer that Mary is Intelligent.

The above example may seem quite trivial, but it is easy to imagine that, with large ontologies, query answering may be a very complex task. The use of DL reasoners allows OWL ontology applications to answer complex queries, and to provide guarantees about the correctness of the result. This is particularly important if ontology based systems are to be used as components in larger applications, such as the Semantic Web, where the correct functioning of automated processes may depend on their being able to (correctly) answer such queries.

5. ONTOLOGY APPLICATIONS

The availability of tools and reasoning systems such as those mentioned in Section 4 has contributed to the increasingly widespread use of OWL, not only in the Semantic Web per se, but as a popular language for ontology development in fields as diverse as biology [28], medicine [9], geography [10], geology [32], astronomy [7], agriculture [30] and defence [23]. Applications of OWL are particularly prevalent in the life sciences where it has been used by the developers of several large biomedical ontologies, including the Biological Pathways Exchange (BioPAX) ontology [27], the GALEN ontology [26], the Foundational Model of Anatomy (FMA) [9], and the National Cancer Institute thesaurus [12].

The importance of reasoning support in such applications was highlighted in [21], which describes a project in which the Medical Entities Dictionary (MED), a large ontology (100,210 classes and 261 properties) that is used at the Columbia Presbyterian Medical Center, was converted into OWL, and checked using an OWL reasoner. This check revealed “systematic modelling errors”, and a significant number of missed subclass relationships which, if not corrected,

“could have cost the hospital many missing results in various decision support and infection control systems that routinely use MED to screen patients”.

6. FUTURE DIRECTIONS

As we have seen in Section 5, OWL is already being successfully used in many applications. This success brings with it, however, many challenges for the future development of both the OWL language and OWL tool support. Central to these is the familiar tension between the requirements for advanced features, in particular increased expressive power, and raw performance, in particular the ability to deal with very large ontologies and data sets.

Use of OWL in the life sciences domain has brought to the fore examples of both of the above mentioned requirements. On the one hand, ontologies describing complex systems in medicine and biology often require expressive power beyond what is currently supported in OWL. Two particular features that are very often requested are the ability to “qualify” cardinality constraints, e.g., to describe the hand as having four parts *that are fingers* and one part *that is a thumb*, and the ability to have some characteristics be transferred across transitive part-whole relations, e.g., to capture the fact that a disease affecting a part of an organ affects the organ as a whole. The former feature (so called qualified cardinality restrictions) has long been well understood, and has been available for some time in DL reasoners; the latter feature is now also well understood, thanks to recent theoretical work in the DL community [16, 13], and has recently been implemented in DL reasoners.

This happy coincidence of user requirements and extensions in the underlying DLs and reasoning systems has led to a proposal to extend OWL with these and other useful features that have been requested by users, for which effective reasoning algorithms are now available, and that OWL tool developers are willing to support. In addition to those mentioned above, the new features include extra syntactic sugar, extended datatype support, simple metamodelling, and extended annotations. The extended language, called OWL 1.1, is now a W3C member submission (see <http://www.w3.org/Submission/2006/10/>), and is already supported by tools such as Swoop, Protégé and TopBraid Composer.

As well as increased expressive power, applications may also bring with them requirements for scalability that are a challenge to current systems. This may include the ability to reason with very large ontologies, perhaps containing 10s or even 100s of thousands of classes, and the ability to use an ontology with very large data sets, perhaps containing 10s or even 100s of millions of individuals—in fact data sets much larger than this will certainly be a requirement in some applications. Researchers are rising to these challenges by developing new reasoning systems such as the OWL Instance Store [2], that uses a combination of DL reasoning and relational database systems to deal with large volumes of instance data, Hermit (see <http://www.cs.man.ac.uk/~bmotik/Hermit/>), that uses a hypertext based technique to deal more effectively with large and complex ontologies, and Kaon2 [18], that reduces OWL ontologies to disjunctive datalog programs, and uses deductive database techniques to enable it to deal with very large data sets.

7. CONCLUSIONS

As we have seen, the goal of Semantic Web research is to transform the Web from a linked document repository into a distributed knowledge base and application platform, thus allowing the vast range of available information and services to be more effectively exploited. As a first step in this transformation, languages such as OWL have been developed; these languages are designed to capture the knowledge that will enable applications to better understand Web accessible resources, and to use them more intelligently. It is easy to imagine that these developments might also be of benefit to users with cognitive or sensory impairments.

Although realising the Semantic Web still seems some way off, OWL has already been very successful, and has rapidly become a de facto standard for ontology development in fields as diverse as geography, geology, astronomy, agriculture, defence and the life sciences. An important factor in this success has been the availability of sophisticated tools with built in reasoning support.

The use of OWL in large scale applications has brought with it new challenges, both with respect to expressive power and scalability, but recent research has also shown how the OWL language and OWL tools can be extended and adapted to meet these challenges.

8. REFERENCES

- [1] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2003.
- [2] S. Bechhofer, I. Horrocks, and D. Turi. The OWL instance store: System description. In *Proc. of the 20th Int. Conf. on Automated Deduction (CADE-20)*, Lecture Notes in Artificial Intelligence, pages 177–181. Springer, 2005.
- [3] T. Berners-Lee. Semantic web road map, Sept. 1998. Available at <http://www.w3.org/DesignIssues/Semantic.html>.
- [4] R. J. Brachman and J. G. Schmolze. An overview of the KL-ONE knowledge representation system. *Cognitive Science*, 9(2):171–216, Apr. 1985.
- [5] B. Cuenca Grau, I. Horrocks, Y. Kazakov, and U. Sattler. Just the right amount: Extracting modules from ontologies. In *Proc. of the Sixteenth International World Wide Web Conference (WWW 2007)*, 2007.
- [6] B. Cuenca Grau, Y. Kazakov, I. Horrocks, and U. Sattler. A logical framework for modular integration of ontologies. In *Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 2007)*, 2007.
- [7] S. Derriere, A. Richard, and A. Preite-Martinez. An ontology of astronomical object types for the virtual observatory. *Proc. of Special Session 3 of the 26th meeting of the IAU: Virtual Observatory in Action: New Science, New Technology, and Next Generation Facilities*, 2006.
- [8] D. Fensel, F. van Harmelen, I. Horrocks, D. McGuinness, and P. F. Patel-Schneider. OIL: An ontology infrastructure for the semantic web. *IEEE Intelligent Systems*, 16(2):38–45, 2001.
- [9] C. Golbreich, S. Zhang, and O. Bodenreider. The foundational model of anatomy in OWL: Experience

- and perspectives. *J. of Web Semantics*, 4(3), 2006.
- [10] J. Goodwin. Experiences of using OWL at the ordnance survey. In *Proc. of the First OWL Experiences and Directions Workshop*, volume 188 of *CEUR Workshop Proceedings*. CEUR (<http://ceur-ws.org/>), 2005.
- [11] V. Haarslev and R. Möller. RACER system description. In *Proc. of the Int. Joint Conf. on Automated Reasoning (IJCAR 2001)*, volume 2083 of *Lecture Notes in Artificial Intelligence*, pages 701–705. Springer, 2001.
- [12] F. W. Hartel, S. de Coronado, R. Dionne, G. Frago, and J. Golbeck. Modeling a description logic vocabulary for cancer research. *Journal of Biomedical Informatics*, 38(2):114–129, 2005.
- [13] I. Horrocks, O. Kutz, and U. Sattler. The even more irresistible *SROIQ*. In *Proc. of the 10th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2006)*, pages 57–67. AAAI Press, 2006.
- [14] I. Horrocks, P. F. Patel-Schneider, and F. van Harmelen. Reviewing the design of DAML+OIL: An ontology language for the semantic web. In *Proc. of the 18th Nat. Conf. on Artificial Intelligence (AAAI 2002)*, pages 792–797. AAAI Press, 2002.
- [15] I. Horrocks, P. F. Patel-Schneider, and F. van Harmelen. From *SHIQ* and RDF to OWL: The making of a web ontology language. *J. of Web Semantics*, 1(1):7–26, 2003.
- [16] I. Horrocks and U. Sattler. Decidability of *SHIQ* with complex role inclusion axioms. *Artificial Intelligence*, 160(1–2):79–104, Dec. 2004.
- [17] I. Horrocks and U. Sattler. A tableaux decision procedure for *SHOIQ*. In *Proc. of the 19th Int. Joint Conf. on Artificial Intelligence (IJCAI 2005)*, pages 448–453, 2005.
- [18] U. Hustadt, B. Motik, and U. Sattler. Reducing SHIQ-description logic to disjunctive datalog programs. In *Proc. of the 9th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2004)*, pages 152–162, 2004.
- [19] A. Kalyanpur, B. Parsia, E. Sirin, B. Cuenca-Grau, and J. Hendler. SWOOP: a web ontology editing browser. *J. of Web Semantics*, 4(2), 2005.
- [20] A. Kalyanpur, B. Parsia, E. Sirin, and J. Hendler. Debugging unsatisfiable classes in owl ontologies. *J. of Web Semantics*, 3(4):243–366, 2005.
- [21] A. Kershenbaum, A. Fokoue, C. Patel, C. Welty, E. Schonberg, J. Cimino, L. Ma, K. Srinivas, R. Schloss, and J. W. Murdock. A view of OWL from the field: Use cases and experiences. In *Proc. of the Second OWL Experiences and Directions Workshop*, volume 216 of *CEUR* (<http://ceur-ws.org/>), 2006.
- [22] H. Knublauch, R. Ferguson, N. Noy, and M. Musen. The Protégé OWL Plugin: An open development environment for semantic web applications. In S. A. McIlraith, D. Plexousakis, and F. van Harmelen, editors, *Proc. of the 2004 International Semantic Web Conference (ISWC 2004)*, number 3298 in *Lecture Notes in Computer Science*, pages 229–243. Springer, 2004.
- [23] L. Lacy, G. Aviles, K. Fraser, W. Gerber, A. Mulvehill, and R. Gaskill. Experiences using OWL in military applications. In *Proc. of the First OWL Experiences and Directions Workshop*, volume 188 of *CEUR Workshop Proceedings*. CEUR (<http://ceur-ws.org/>), 2005.
- [24] S. McIlraith, T. C. Son, and H. Zeng. Semantic web services. *IEEE Intelligent Systems*, 16:46–53, 2001.
- [25] P. F. Patel-Schneider, P. Hayes, and I. Horrocks. OWL Web Ontology Language semantics and abstract syntax. W3C Recommendation, 10 February 2004. Available at <http://www.w3.org/TR/owl-semantics/>.
- [26] A. Rector and J. Rogers. Ontological and practical issues in using a description logic to represent medical concept systems: Experience from GALEN. In *Reasoning Web, Second International Summer School, Tutorial Lectures*, volume 4126 of *LNCS*, pages 197–231. SV, 2006.
- [27] A. Ruttenberg, J. Rees, and J. Luciano. Experience using OWL DL for the exchange of biological pathway information. In *Proc. of the First OWL Experiences and Directions Workshop*, volume 188 of *CEUR Workshop Proceedings*. CEUR (<http://ceur-ws.org/>), 2005.
- [28] A. Sidhu, T. Dillon, E. Chang, and B. S. Sidhu. Protein ontology development using OWL. In *Proc. of the First OWL Experiences and Directions Workshop*, volume 188 of *CEUR Workshop Proceedings*. CEUR (<http://ceur-ws.org/>), 2005.
- [29] E. Sirin, B. Parsia, B. Cuenca Grau, A. Kalyanpur, and Y. Katz. Pellet: A practical owl-dl reasoner. To appear, 2005.
- [30] D. Soergel, B. Lauser, A. Liang, F. Fisseha, J. Keizer, and S. Katz. Reengineering thesauri for new applications: The AGROVOC example. *J. of Digital Information*, 4(4), 2004.
- [31] K. Spackman. Managing clinical terminology hierarchies using algorithmic calculation of subsumption: Experience with SNOMED-RT. *J. of the Amer. Med. Informatics Ass.*, 2000. Fall Symposium Special Issue.
- [32] Semantic web for earth and environmental terminology (SWEET). Jet Propulsion Laboratory, California Institute of Technology, 2006. <http://sweet.jpl.nasa.gov/>.
- [33] D. Tsarkov and I. Horrocks. FaCT++ description logic reasoner: System description. In *Proc. of the Int. Joint Conf. on Automated Reasoning (IJCAR 2006)*, volume 4130 of *Lecture Notes in Artificial Intelligence*, pages 292–297. Springer, 2006.
- [34] R. Volz, S. Handschuh, S. Staab, L. Stojanovic, and N. Stojanovic. Unveiling the hidden bride: Deep Annotation for Mapping and Migrating Legacy Data to the Semantic Web. *Journal of Web Semantics*, 2004.
- [35] W. A. Woods. What’s in a link: Foundations for semantic networks. In R. J. Brachman and H. J. Levesque, editors, *Readings in Knowledge Representation*, pages 217–241. Morgan Kaufmann Publishers, San Francisco, California, 1985.