

# The OWL API: A Java API for OWL Ontologies

Matthew Horridge<sup>a</sup>, Sean Bechhofer<sup>a,\*</sup>

<sup>a</sup> *School of Computer Science, Kilburn Building, Oxford Road, University of Manchester, M13 9PL, United Kingdom*  
email: *first.second@manchester.ac.uk*

**Abstract.** We present the OWL API, a high level Application Programming Interface (API) for working with OWL ontologies. The OWL API is closely aligned with the OWL 2 structural specification. It supports parsing and rendering in the syntaxes defined in the W3C specification (Functional Syntax, RDF/XML, OWL/XML and the Manchester OWL Syntax); manipulation of ontological structures; and the use of reasoning engines. The reference implementation of the OWL API, written in Java, includes validators for the various OWL 2 profiles - OWL 2 QL, OWL 2 EL and OWL 2 RL. The OWL API has widespread usage in a variety of tools and applications.

Keywords: OWL, API, Java, reasoning, application development

## 1. Introduction

The Web Ontology Language OWL has been a W3C Recommendation [32] since 2004, with OWL 2 [41] refining and extending the original specification. A pre-requisite for the success and adoption of languages is the existence of tooling, and OWL is no exception, with a wide range of (free, open-source and commercial) tools now available. These tools support the creation and editing of OWL ontologies, reasoning over ontologies, and the use of ontologies in applications.

In this paper, we discuss the OWL API, a high level Application Programming Interface (API) that supports the creation and manipulation of OWL Ontologies. The OWL API has been available since 2003, and has undergone a number of design revisions, in particular tracking the evolution of OWL itself. Key aspects of the OWL API include an axiom-centric abstraction, first class change support, general purpose reasoner interfaces, validators for the various OWL 2 profiles, and support for parsing and serialising ontologies in a variety of syntaxes. The OWL API also has a flexible

design that allows third parties to provide alternative implementations for all major components. The OWL API is implemented in Java and is available as open source under an LGPL licence<sup>1</sup>.

Since its initial development in 2003 [8], the OWL API has been used in multiple development projects, including Protégé-4 [25], SWOOP [23], the NeOn Toolkit [17]<sup>2</sup>, OWLSight<sup>3</sup>, OntoTrack [26], the Pellet reasoner [36]<sup>4</sup> and a number of online validation and conversion services<sup>5</sup>. Work on early versions of the OWL API contributed valuable implementation experience [3] to the WebOnt Working Group<sup>6</sup>, facilitating the production of the OWL Recommendation. Across all versions, the OWL API has over 34,000 downloads (with 500+ downloads per month over the last year), reinforcing our view that it has played, and continues to play, a key role in promoting the uptake and use of OWL.

---

<sup>1</sup><http://owlapi.sourceforge.net>

<sup>2</sup><http://www.neon-toolkit.org/>

<sup>3</sup><http://pellet.owldl.com/ontology-browser/>

<sup>4</sup><http://clarkparsia.com/pellet>

<sup>5</sup><http://owl.cs.manchester.ac.uk>

<sup>6</sup><http://www.w3.org/2001/sw/WebOnt>

---

\*Corresponding author

## 2. OWL API Design Philosophy

The original inspiration for the OWL API came from observations of the impact made by the provision of the XML Document Object Model (DOM) [40]. The DOM, along with freely available implementations (such as the Java implementations in Sun’s JDK [38]) allowed a large number of developers to use and manipulate XML in applications, which turn facilitated the widespread adoption of XML. Our belief was that a similar effort would prove of great benefit to the adoption of OWL.

The provision of an OWL API can allow developers to work at an appropriate level of abstractions, isolating them from potential issues related to, for example, serialisation and parsing of data structures. This was a particular consideration with OWL, given OWL’s close relationship with RDF [24] and its triple-based representation.

In order to provide this “high level view”, the design of the OWL API is directly based on the OWL 2 Structural Specification [30]. An ontology is simply viewed as a set of axioms and annotations as depicted in Figure 1. The names and hierarchies of interfaces for entities, class expressions and axioms in the OWL API correspond closely to the structural specification, relating the high level OWL 2 specification directly to the design of the OWL API. The OWL API supports loading and saving ontologies in a variety of syntaxes. However, none of the model interfaces in the OWL API reflect, or are biased to any particular concrete syntax or model. For example, unlike other APIs such as Jena [12], or the Protégé 3.X API, the representation of class expressions and axioms is *not* at the level of RDF triples.

## 3. Detailed Design Principles

Although aspects of the API design have evolved since its original inception and implementation, a number of basic design principles have endured. These are outlined below, but can be summarised as:

- Read-only model interfaces with explicit change operations.
- Independence from concrete serialisations (in particular triple based representations).
- A clear separation between components providing particular functionalities, such as representation, manipulation, parsing, rendering.
- Separation of assertion and inference.

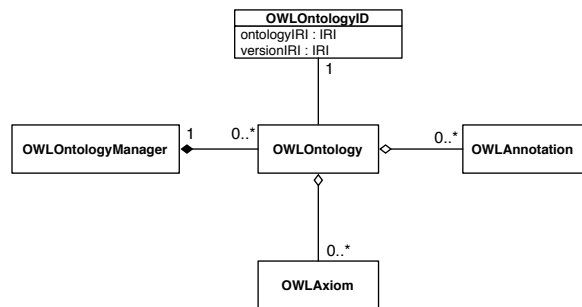


Fig. 1. A UML diagram showing the management of ontologies in the OWL API. The `OWLOntology` interface provides accessors for efficiently obtaining the axioms contained within an ontology. The method of storing axioms is provided by different implementations of the `OWLOntology` interface. The design of the API makes it possible to mix and match implementations, for example, an in-memory ontology could be used together with one that is stored in a database and one that is stored in some kind of triple store. The OWL API reference implementation provides an efficient in-memory storage solution.

### 3.1. Model

The OWL API’s model provides access to an OWL ontology through a number of Interfaces and Class definitions. The model interfaces are *read-only*, in that they do not provide explicit functionality for change of the underlying data structures (see Section 3.2).

Following the OWL 2 Structural Specification, the model provides an *axiom-centric* view on OWL ontologies (see Figure 1), with an `OWLOntology` containing a number of `OWLXiom` objects. Convenience methods are also available that will answer questions of containment such as “does Ontology O contain class C?”, where containment is taken to mean that the ontology contains an axiom referring to C. This supports what we might call *locality of information*, in that all assertions about classes are associated with a particular ontology. As OWL operates in an open web context, different ontologies can make different assertions about the same classes and properties.

Although this “axiom-centric” view differs from the original OWL API design, which provided a “frame-oriented” model, the desire to provide a level of abstraction that insulates developers and applications from underlying syntactic presentations, in particular RDF or triple-based representations has been a cornerstone of the OWL API’s approach.

The model data structures make extensive use of the Visitor pattern [14] which allows separation of data structure from functionality. The Visitor makes it easier to add functionality to a class hierarchy (and is par-

ticularly suited to situations where data structures represent abstract syntax). The Visitor pattern does, however, have the drawback that changes to the data structure can be expensive as it may require changes to the Visitor implementations. For the OWL API, however, where the data structures are stable, the pattern is a good fit.

The specification of the data model is largely through Java Interfaces, enabling the use of alternative implementation strategies. The reference implementation provides data structures for efficient in-memory representations of ontologies. For many purposes this is sufficient. For example, recent versions of the National Cancer Institute (NCI) ontology can comfortably fit into around 700MB of memory (100MB of memory without annotations). However, the API has been designed so that it is possible to provide other storage mechanisms for ontologies such as relational databases or triple stores and to “mix and match” storage implementations, so that an ontology imports closure could contain in-memory representations of ontologies, ontologies persisted in secondary storage in the form of a custom database, and ontologies stored in triple stores.

While the API does not include these alternative storage mechanisms out of the box, third parties have developed such solutions. An exemplar solution, called OWLDB, has been developed by Hens et al [20]. Their solution stores ontologies in a relational database, using a “native” mapping of OWL constructs (as opposed to a mapping to triples) to a custom database schema. A similar approach is taken by Redmond in [34].

### 3.1.1. Ontology Management

Besides the model interfaces for representing entities, class expressions and axioms, it is necessary to allow applications to manage ontologies. Figure 1 shows a high level overview of this picture. The `OWLontology` interface provides a point for accessing the axioms contained in an ontology. As discussed above, different implementations of the `OWLontology` interface can provide different storage mechanisms for ontologies.

The `OWLontologyManager` provides a central point for creating, loading, changing and saving ontologies, which are instances of the `OWLontology` interface. Each ontology is created or loaded by an ontology manager. Each instance of an ontology is unique to a particular manager, and all changes to an ontology are applied via its manager.

This centralised management design allows client applications to have a single access point to ontologies, to provide redirection mechanisms and other customisations for loading ontologies, and allows client applications to monitor all changes that are applied to any loaded ontologies. The manager also hides much of the complexity associated with choosing the appropriate parsers and renderers for loading and saving ontologies.

### 3.2. Change

The core Ontology model interfaces are *read-only*, with issues of change (additions and removals of axioms) handled through the use of explicit change objects. Class `OWLontologyChange` provides a top level change object with further subclasses defined that encapsulate particular changes. Change enactment is then provided through a further use of the Visitor pattern, with `OWLontologyChangeVisitor` objects enacting changes.

This use of explicit change objects follows the Command design pattern [14] which encapsulates a change request as an object. Changes are then enacted by a Visitor object.

This use of the Command pattern facilitates support for operations such as undo or redo, and the encapsulation of changes as operations provides a mechanism with which to track changes and support version management. Change objects also provide a convenient place for storing metadata about the changes, for example the user who requested the change – information which is again crucial in supporting the ontology management and editing process.

All ontology changes are applied through an ontology manager. This means that it is possible to apply a list of ontology changes, which make multiple changes to multiple ontologies, as a single unit. This works well for applications such as ontology editors, where an edit operation such as entity name change (entity IRI change) can involve the addition and removal of multiple axioms from multiple ontologies—it is possible to group the changes together to form one “editor operation” and apply these changes at one time.

### 3.3. Inference and Reasoning

A key aspect of OWL is the provision of a semantics that defines precisely what entailment means with respect to OWL ontologies, and provides formal descriptions of properties such as consistency. The im-

plementation of these semantics is a non-trivial matter. By separating reasoning functionality, we can relieve implementors of the burden of this, while allowing those who do provide such implementations to be explicit about this in their advertised functionality. In addition, as is discussed in some detail in [8], separation of assertion and inference is important in the implementation of OWL, particularly when developing user applications, as users may need explicit indication as to why, for example, hierarchical relationships are present.

The `OWLReasoner` interface provides access to functionality relating to the process of *reasoning* with OWL ontologies, supporting tasks such as consistency checking, computation of class or property hierarchies and axiom entailment. Of course, providing method signatures does not go all the way to advertising the functionality of an implementation – there is no guarantee that a component implementing the interface necessarily implements the semantics correctly. However, signatures and extensive detailed documentation of reasoner interfaces go some way towards providing an expectation of the operations that are being supported. From a client perspective, the benefits of having a well defined reasoner interface cannot be overstated. The ability to exchange one reasoner implementation for another in a client application, and know that, in terms of query answering, each reasoner implementation should exhibit the same behaviour is a major advantage. Collections of test data (such as the OWL Test Cases [11]) can allow systematic testing and a level of confidence as to whether the implementation is, in fact, performing correctly.

The reasoner interfaces have been designed so as to enable reasoners to expose functionality that provides *incremental* reasoning support. The OWL API allows a reasoner to be initialised so that it listens for ontology changes and either immediately processes the changes or queues them in a buffer which can later be processed. This design caters for the scenario where a reasoner is used within an ontology editor and, at some point, must respond when prompted to edits of the ontology, or situations where a reasoner should respond to ontology changes as they arrive.

A number of existing implementations including the CEL [1], FaCT++ [39], HermiT [29], Pellet [36], and Racer Pro [16] and SnoRocket reasoners provide OWL API wrappers. These reasoning engines are thus easily integrated into OWL API based applications such as Protégé-4 or the NeOn Toolkit.

### 3.4. Parsing and Rendering OWL Ontologies

A benefit of aligning the OWL API with the OWL 2 structural specification is that there is no commitment to a particular concrete syntax. Conforming [37] OWL implementations or tools *must* support RDF/XML, but there are several other syntaxes that exist which are optimised for different purposes. For example, Turtle syntax [9] provides a “compact and natural text form” of RDF serialisation. The Manchester OWL Syntax [18] provides a human readable serialisation for OWL ontologies.

The OWL API includes out of the box support for reading and writing ontologies in several syntaxes, including RDF/XML, Turtle, OWL/XML, OWL Functional Syntax, The Manchester OWL Syntax, KRSS Syntax<sup>7</sup> and the OBO flat file format<sup>8</sup>. Due to the underlying design of the API, it is possible for the imports closure of an ontology to contain ontologies that were parsed from ontology documents written in different syntaxes.

The reference implementation of the OWL API uses a registry of parsers and renderers, supporting the addition of custom syntaxes. The appropriate parser is *automatically selected* at runtime when an ontology is loaded. By default, ontologies are saved back into the format from which they were parsed, but it is possible to override this in order to perform syntax conversion tasks and “save as” operations in editors for example.

## 4. Evolution

Three versions of the OWL API have been developed. The first version [8] drew on earlier experiences in the implementation of OilEd [5], one of the first user tools to use Description Logic reasoning services (through use of the DIG protocol) to assist users in the construction of OIL and DAML+OIL ontologies. Although OilEd was widely used (with over 10,000 download requests), the underlying implementation suffered from problems, not least of which was the lack of a separation between interface and implementation. The first version of the OWL API addressed a number of these concerns, while carrying forward other design decisions such as extensive use of the Visitor pattern [14].

<sup>7</sup><http://dl.kr.org/krss-spec.ps>

<sup>8</sup>[http://www.geneontology.org/GO.format.obo-1\\_2.shtml](http://www.geneontology.org/GO.format.obo-1_2.shtml)



## 5.2. Explanation

In recent years, explanation of entailments has been a “hot topic”. Many of the afore mentioned ontology development environments now feature the ability to generate explanations for unwanted or surprising entailments that arise during editing or browsing ontologies. Many of these explanation facilities are based on justifications [35,2,22], which are minimal sets of axioms that are sufficient for an entailment to hold. The OWL API ships with code for generating these explanations, making it easy to add basic explanation services to client applications.

## 5.3. Modularity

Many applications require the ability to work with well defined portions of an ontology, that are usually called modules. The OWL API contains basic functionality for extracting syntactic-locality based modules [13] from ontologies. These kinds of modules are based on the notion of conservative-extensions [15], and guarantee to preserve entailments from an ontology over a given signature.

## 6. Other Frameworks

There have been a number of similar initiatives to provide application interfaces aimed at OWL. HP’s Jena Toolkit [12] supplies ontology interfaces that provides convenience wrappers around RDF interfaces. Comparisons of the OWL API and Jena’s triple-based approach for some tasks are explored in [4].

The KAON [10] toolkit is an open-source ontology management infrastructure targeted for business applications<sup>11</sup>. The ontology model supported in KAON is less expressive than that described by OWL since an important focus of KAON is performance reasoning on knowledge bases with simple ontologies and large ABoxes [28]. Despite this, many of our underlying design considerations conceptually follow the KAON design.

## 7. Download, Takeup and Usage

The OWL API is available as open source under the LGPL licence and can be downloaded from Source-

forge<sup>12</sup>. Statistics taken from sourceforge shown in Figure 3 indicate a steady increase in downloads and web traffic<sup>13</sup>. Since version 1 was released, the OWL API has had over 34,000 downloads from the Sourceforge website. It should be noted that this count does not include copies of the API that are now distributed with widely used tools such as Protégé-4, the NeOn Toolkit, Pellet and HermiT.

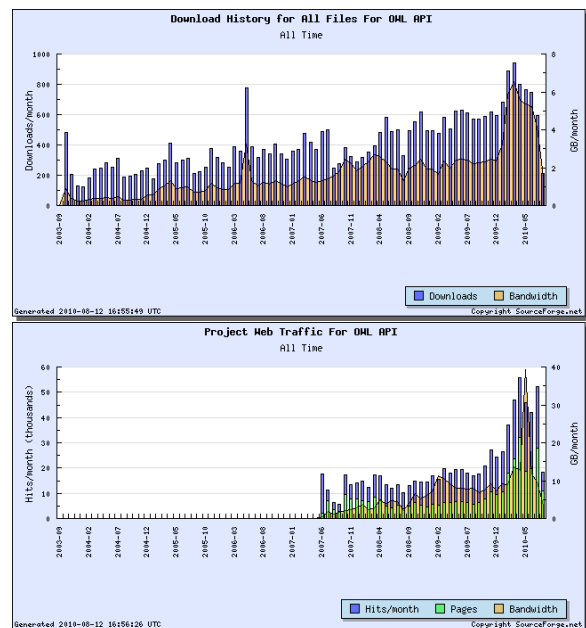


Fig. 3. Sourceforge statistics for downloads and web traffic

Some examples of the kinds of applications and services that have been developed using the OWL API are discussed below. As the codebase is freely available, we fully expect there to be further third-party usage of the OWL API that we are not explicitly aware of.

**Protégé-4** [25] is an open source OWL ontology editor that was initially designed and developed at the University of Manchester. Protégé-4 uses the OWL API to underpin all ontology management tasks, from loading and saving ontologies, to manipulating ontologies during editing, to interacting and offering a choice of OWL reasoners. Virtually all of the functionality provided by the OWL API is utilised by Protégé-4.

The **NeOn Toolkit**<sup>2</sup> [17] is an Eclipse based ontology development environment that was developed as part of the 14.7 million Euro EU funded project.

<sup>12</sup><http://owlapi.sourceforge.net>

<sup>13</sup>Note that the figures for project web traffic are only available from mid- 2007.

<sup>11</sup>Available at <http://kaon.semanticweb.org>

While early versions of the toolkit were written on top of KAON2<sup>14</sup>, in 2010, the project moved to the OWL API<sup>15</sup>.

**OWLSight**<sup>3</sup> is a web based ontology browser written by Clark & Parsia that uses the Pellet reasoner. The browser is written using the Google Web Toolkit, with the OWL API being used to read and access ontologies.

The **Ontology Browser**<sup>16</sup> dynamically generates documentation for ontologies and is based on the OWLDoc software. The OWL API is used for loading and accessing ontologies and interfacing with the FaCT++ reasoner.

**OntoTrack** [26] is a browsing and editing tool for OWL ontologies that is developed at Ulm University. The OWL API is used for loading and accessing ontologies for rendering into a graph.

The **SKOS API**<sup>17</sup> [21] is a Java interface and implementation for SKOS [27], built on top of the OWL API. It adopts a similar approach to the OWL API, providing an abstract data model for SKOS that avoids commitment to particular concrete syntaxes. SKOSEd, a SKOS editor is implemented as a plugin in to Protégé-4.

The **OWLlink API** [31] implements the OWLlink protocol on top of the OWL API. Besides providing an API to access remote OWLlink reasoning engines, it turns any OWL API aware reasoner into an OWLlink server. In essence OWLlink and the OWLlink API replace the DIG [6] protocol.

A number of online services<sup>18</sup> are implemented using the OWL API.

**Syntax Converter** A web based application that converts ontologies written in one OWL syntax to another OWL syntax. Syntax is converted from one format to another format by loading and saving ontologies using the OWL API.

**Ontology Repository** A repository of ontologies used for reasoner and tools testing. The loading and serialising makes use of the OWL API parsers and renderers, while the metrics for each ontology are computed by the OWL API's Metrics API.

**Validator** Uses the various profile validators to determine if an ontology and its imports closure is within

a specific profile. The validator returns validation reports in a variety of human readable syntaxes. The validator makes use of the Profiles API, which is part of the OWL API.

**Module Extractor** Allows locality based modules to be extracted from ontologies. The extractor makes use of OWL API modularisation code, which currently provides Syntactic Locality Based Modules [19].

**Metrics** Allows ontologies to be submitted and returns a report about the number and types of axioms in an ontology and its imports closure. This application uses the Metrics API, which is part of the OWL API.

## 8. Limitations

As discussed above, the OWL API has proved to be a popular and useful code library. The design decisions made in the development of the OWL API do, however, impact on its suitability for application usage. The most crucial of these is that the OWL API is primarily designed to be an application that supports manipulation of OWL ontologies at a particular level of abstraction – which is not the RDF level. Thus the API is perhaps less suitable for those wishing to explicitly exploit the layering of OWL on RDF. In addition, the original design of the OWL API was “tuned” to target the OWL DL species. Although OWL Full ontologies can be manipulated using the OWL API, alternative frameworks such as Jena may be more suitable if OWL Full aspects are to be exploited in an application.

The reference implementation supplied with the OWL API downloads uses an in-memory representation. This places some restriction on the size of ontologies that can be processed using the OWL API. In practice, we have not found this to be a serious issue. Table 1 shows load times (in seconds) and memory consumption (in megabytes) after loading for several well known large ontologies. For comparison, the load times and memory consumption for the Jena API are also given, showing comparable performance.

## 9. Summary

The OWL API provides a collection of powerful and flexible interfaces supporting the use of OWL ontologies within applications. The model explicitly supports the recent OWL 2 Recommendation. Common interfaces to reasoning engines are defined, facilitating the

<sup>14</sup><http://kaon2.semanticweb.org>

<sup>15</sup>[http://www.neon-project.org/nw/NeOn\\_Toolkit\\_-\\_Latest\\_Major\\_Release](http://www.neon-project.org/nw/NeOn_Toolkit_-_Latest_Major_Release)

<sup>16</sup><http://owl.cs.manchester.ac.uk/browser>

<sup>17</sup><http://skosapi.sourceforge.net/>

<sup>18</sup><http://owl.cs.manchester.ac.uk>

Ontology	Axioms	Triples	Load Time	Mem.	
NCI	1,093,733	1,449,558	19.0 s	638 Mb	O/A
			36.3 s	600 Mb	Jen.
FMA	1,651,533	1,686,457	26.5 s	831 Mb	O/A
			43.8 s	482 Mb	Jen.
SNOMED	1,036,800	3,355,279	30.1 s	770 Mb	O/A
			59.4 s	875 Mb	Jen.

Table 1

Load times and memory usage for well known large ontologies. The OWL API (O/A) is compared with Jena (Je.)

use of inference within applications. The distribution includes a reference, main-memory implementation, and provides support for parsing and rendering ontologies in a variety of concrete syntaxes along with OWL 2 Profile validators. The API is in widespread usage, with a growing user community.

The overview presented here is necessarily brief – space constraints prevent us from providing in depth descriptions of the design and details of the code base. Further detailed documentation is available from the OWL API Sourceforge site<sup>12</sup> along with tutorial presentations, examples and full javadoc documentation for the code base. There is also an active mailing list for developers<sup>19</sup>.

## 10. Acknowledgements

The OWL API has required extensive programming and design effort, and the level of sophistication and robustness achieved would have been difficult, if not impossible, without key contributions from a number of people. The authors would like to acknowledge those who have contributed to the OWL API since the first version was designed and developed by Sean Bechhofer, Raphael Volz and Philip Lord in 2003 [7]. In particular, those who have made significant contributions (in alphabetical order) include: Ron Alford<sup>20</sup>, Nick Drummond<sup>21</sup>, Birte Glimm<sup>22</sup> (reasoner interface design), Luigi Iannone<sup>21</sup>, Aditya Kalyanpur<sup>23</sup>, Pavel Klinov<sup>21</sup>, Boris Motik<sup>22</sup> (reasoner interface and general design), Olaf Noppens<sup>24</sup> (reasoner interface design and parsing contributions), Ignazio Palmisano<sup>21</sup>

<sup>19</sup>[http://sourceforge.net/mailarchive/forum.php?forum\\_name=owlapi-developer](http://sourceforge.net/mailarchive/forum.php?forum_name=owlapi-developer)

<sup>20</sup>Clark & Parsia

<sup>21</sup>The University of Manchester

<sup>22</sup>Oxford University

<sup>23</sup>IBM TJ Watson Research Center, New York

<sup>24</sup>Ulm University

(concurrency support), Timothy Redmond<sup>25</sup>, Angus Roberts<sup>21</sup>, Nicolas Rouquette<sup>26</sup>, Thomas Schneider<sup>21</sup> (modularity support), Evren Sirin<sup>20</sup> (reasoner interface design, explanation code), Mike Smith<sup>20</sup>, and Daniele Turi<sup>21</sup>

Thanks also go to past and present members of the Information Management Group (IMG) and the Bio-Health Informatics Group (BHIG) at the University of Manchester, and all of the OWL API users who have posted bug reports and feedback on the OWL API public mailing list.

Development of the OWL API has been funded from a number of sources, including the EU funded Wonderweb project<sup>27</sup>; the JISC funded CO-ODE project<sup>28</sup> and the EU funded TONES project<sup>29</sup>.

## References

- [1] F. Baader, C. Lutz, and B. Suntisrivaraporn. CEL—a polynomial-time reasoner for life science ontologies. In U. Furbach and N. Shankar, editors, *Proceedings of the 3rd International Joint Conference on Automated Reasoning (IJ-CAR'06)*, volume 4130 of *Lecture Notes in Artificial Intelligence*, pages 287–291. Springer-Verlag, 2006.
- [2] Franz Baader and Bernhard Hollunder. Embedding defaults into terminological knowledge representation formalisms. *Journal of Automated Reasoning*, 14(1):149–180, 1995.
- [3] Sean Bechhofer. OWL Web Ontology Language Parsing OWL in RDF/XML. W3C Working Group Note, World Wide Web Consortium, January 2004. <http://www.w3.org/TR/owl-parsing>.
- [4] Sean Bechhofer and Jeremy J. Carroll. Parsing OWL DL: Trees or Triples? In *Proceedings of the World Wide Web Conference, WWW2004*, pages 266–275. ACM Press, 2004.
- [5] Sean Bechhofer, Ian Horrocks, Carole Goble, and Robert Stevens. OilEd: a Reason-able Ontology Editor for the Semantic Web. In *Proceedings of KI2001, Joint German/Austrian conference on Artificial Intelligence*, volume 2174 of *LNAI*, pages 396–408. Springer-Verlage, Sep 2001.
- [6] Sean Bechhofer, Ralf Möller, and Peter Crowther. The DIG Description Logic Interface. In *Proceedings of DL2003 International Workshop on Description Logics*, September 2003.
- [7] Sean Bechhofer, Raphael Volz, and Philip Lord. Cooking the semantic web with the OWL API. In Dieter Fensel, Katia Sycara, and John Mylopoulos, editors, *The Semantic Web - ISWC 2003. The Second International Semantic Web Conference, Sanibel Island, Florida, USA*, volume 2870/2003 of *Lecture Notes in Computer Science*, pages 659–675, Sanibel Island, Florida, USA, October 2003. Springer.

<sup>25</sup>Stanford University

<sup>26</sup>JPL, NASA

<sup>27</sup><http://wonderweb.semanticweb.org/> EU IST-2201-33052

<sup>28</sup><http://www.co-ode.org>

<sup>29</sup><http://www.tonesproject.org/> FP6-7603



- [8] Sean Bechhofer, Raphael Volz, and Phillip Lord. Cooking the Semantic Web with the OWL API. In Dieter Fensel, Katia Sycara, and John Mylopoulos, editors, *2nd International Semantic Web Conference, ISWC*, volume 2870 of *Lecture Notes in Computer Science*. Springer, October 2003.
- [9] David Beckett and Tim Berners-Lee. Turtle - Terse RDF Triple Language. Team Submission, World Wide Web Consortium, 2008. <http://www.w3.org/TeamSubmission/turtle/>.
- [10] Erol Bozsak, Marc Ehrig, Siegfried Handschuh, Andreas Hotho, Alexander Maedche, Boris Motik, Daniel Oberle, Christoph Schmitz, Steffen Staab, Ljiljana Stojanovic, Nenad Stojanovic, Rudi Studer, Gerd Stumme, York Sure, Julien Tane, Raphael Volz, and Valentin Zacharias. KAON - Towards a Large Scale Semantic Web. In K. Bauknecht, A. Min Tjoa, and G. Quirchmayr, editors, *EC-Web 2002*, volume 2455 of *Lecture Notes in Computer Science*, pages 304–313. Springer, September 2002.
- [11] J. Carroll and J. De Roo. OWL Web Ontology Language Test Cases. W3C Recommendation, World Wide Web Consortium, 2004. <http://www.w3.org/TR/owl-test/>.
- [12] Jeremy J. Carroll, Ian Dickinson, Chris Dollin, Dave Reynolds, Andy Seaborne, and Kevin Wilkinson. Jena: implementing the semantic web recommendations. In Stuart Feldman, Mike Uretsky, Mark Najork, and Craig Wills, editors, *Proceedings of the 13th international World Wide Web conference on Alternate track papers & posters*, pages 74–83, New York, NY, USA, May 2004. ACM.
- [13] Bernardo Cuenca Grau, Ian Horrocks, Yevgeny Kazakov, and Ulrike Sattler. Just the right amount: Extracting modules from ontologies. In *WWW 2007, Proceedings of the 16th International World Wide Web Conference, Banff, Canada, May 8-12, 2007*, pages 717–727, 2007.
- [14] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Professional Computing Series. Addison-Wesley, 1995.
- [15] Silvio Ghilardi, Carsten Lutz, and Frank Wolter. Did i damage my ontology? a case for conservative extensions in description logic. In Patrick Doherty, John Mylopoulos, and Christopher A. Welty, editors, *The 10th International Conference on Principles of Knowledge Representation and Reasoning (KR 2006)*, Lake District, United Kingdom. AAAI Press, June 2006.
- [16] Volker Haarslev and Ralf Möller. RACER system description. In *International Joint Conference on Automated Reasoning (IJCAR 2001)*, volume 2083 of *Lecture Notes In Computer Science*, pages 701–705, 2001.
- [17] Peter Haase, Holger Lewen, Rudi Studer, Duc Thanh Tran, Michael Erdmann, Mathieu d’Aquin, and Enrico Motta. The neon ontology engineering toolkit. In Jeff Korn, editor, *WWW 2008 Developers Track*, April 2008.
- [18] Matthew Horridge and Peter F. Patel-Schneider. OWL 2 Web Ontology Language Manchester Syntax. W3C Working Group Note, World Wide Web Consortium, 2009. <http://www.w3.org/TR/owl2-manchester-syntax/>.
- [19] Ernesto Jiménez-Ruiz, Bernardo Cuenca Grau, Ulrike Sattler, Thomas Schneider, and Raphael Berlanga Llavori. Safe and economic re-use of ontologies: A logic-based methodology and tool support. In Sean Bechhofer, Manfred Hauswirth, Jörg Hoffmann, and Manolis Koubarakis, editors, *The Semantic Web: Research and Applications, 5th European Semantic Web Conference, ESWC 2008, Tenerife, Canary Islands, Spain*, volume 5021 of *Lecture Notes in Computer Science*, pages 185–199. Springer, June 2008.
- [20] Joachim Kleb Jörg Henss and Stephan Grimm. A database backend for OWL. In Rinke Hoeksta and Peter F. Patel-Schneider, editors, *OWL: Experiences and Directions (OWLED 2009)*, CEUR Workshop Proceedings. CEUR-WS.org, October 2009.
- [21] Simon Jupp, Sean Bechhofer, and Robert Stevens. A flexible API and editor for SKOS. In *Proceedings of the 6th European Semantic Web Conference (ESWC), Heraklion, Crete*, volume 5554 of *Lecture Notes in Computer Science*, pages 506–520. Springer, 2009.
- [22] Aditya Kalyanpur. *Debugging and Repair of OWL Ontologies*. PhD thesis, The Graduate School of the University of Maryland, 2006.
- [23] Aditya Kalyanpur, Bijan Parsia, Evren Sirin, Bernardo Cuenca Grau, and James Hendler. Swoop: A web ontology editing browser. *Web Semantics: Science, Services and Agents on the World Wide Web*, 4(2):144 – 153, 2006. Semantic Grid –The Convergence of Technologies.
- [24] Graham Klyne and Jeremy J. Carroll. Resource Description Framework (RDF): Concepts and Abstract Syntax. W3C Recommendation, World Wide Web Consortium, 2004. <http://www.w3.org/TR/owl-guide/>.
- [25] Holger Knublauch, Ray W. Ferguson, Natalya F. Noy, and Mark A. Musen. The protégé owl plugin: An open development environment for semantic web applications. In Sheila A. McIlraith, Dimitris Plexousakis, and Frank van Harmelen, editors, *The Semantic Web – ISWC 2004*, volume 3298 of *Lecture Notes in Computer Science*, pages 229–243. Springer Berlin / Heidelberg, 2004.
- [26] Thorsten Liebig and Olaf Noppens. Ontotrack: Combining browsing and editing with reasoning and explaining for owl lite ontologies. In Sheila A. McIlraith, Dimitris Plexousakis, and Frank van Harmelen, editors, *he Semantic Web – ISWC 2004*, volume 3298 of *Lecture Notes in Computer Science*, pages 244–258. Springer Berlin / Heidelberg, 2004.
- [27] Alistair Miles and Sean Bechhofer. SKOS Simple Knowledge Organization System Reference. W3C Recommendation, World Wide Web Consortium, 2009. <http://www.w3.org/TR/skos-reference/>.
- [28] Boris Motik, Alexander Maedche, and Raphael Volz. A conceptual modeling approach for building semantics-driven enterprise applications. In *Proc. of the International Conference on Ontologies, Databases and Applications of SEMantics ODBASE 2002*. Springer, LNCS, 2002.
- [29] Boris Motik, Rob Shearer, and Ian Horrocks. Optimized reasoning in description logics using hypertableaux. In *Proc. of the 21st Int. Conf. on Automated Deduction (CADE-21)*, volume 4603 of *Lecture Notes in Artificial Intelligence*, pages 67–83. Springer, 2007.
- [30] Motik, Boris and Patel-Schneider, Peter F. and Parsia, Bijan. OWL 2 Web Ontology Language Structural Specification and Functional-Style Syntax. W3C Recommendation, World Wide Web Consortium, 2009. <http://www.w3.org/TR/owl2-syntax/>.
- [31] Olaf Noppens, Marko Luther, and Thorsten Liebig. The OWLlink api teaching owl components a common protocol. In Evren Sirin, editor, *OWL: Experiences and Directions*

- (*OWLED 2010*), 2010.
- [32] Peter F. Patel-Schneider, Patrick Hayes, and Ian Horrocks. OWL Web Ontology Language semantics and abstract syntax. W3C Recommendation, World Wide Web Consortium, 10 February 2004.
- [33] Peter F. Patel-Schneider and Ian Horrocks. OWL 1.1 Web Ontology Language Overview. Member Submission, World Wide Web Consortium, 2006. <http://www.w3.org/Submission/owl1-overview/>.
- [34] Timothy Redmond. An open source database backend for the OWL API and protégé 4. In Evren Sirin, editor, *OWL: Experiences and Directions (OWLED 2010)*, 2010.
- [35] Stefan Schlobach and Ronald Cornet. Non-standard reasoning services for the debugging of description logic terminologies. In *IJCAI International Joint Conference on Artificial Intelligence*, 2003.
- [36] Evren Sirin, Bijan Parsia, Bernardo Cuenca Grau, Aditya Kalyanpur, and Yarden Katz. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics*, 5(2), 2007.
- [37] Michael Smith, Ian Horrocks, Markus Krötzsch, and Birte Glimm. OWL 2 Web Ontology Language Conformance. W3C Recommendation, World Wide Web Consortium, 2009. <http://www.w3.org/TR/owl2-conformance/>.
- [38] Sun Microsystems, Inc. Java™Platform. <http://java.sun.com/j2se/>.
- [39] Dmitry Tsarkov and Ian 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.
- [40] W3C DOM Working Group. Document Object Model. <http://www.w3.org/DOM/>.
- [41] W3C OWL Working Group. OWL 2 Web Ontology Language Document Overview. W3C Recommendation, World Wide Web Consortium, 2009. <http://www.w3.org/TR/owl2-overview/>.