

INTEGRATING HIERARCHICAL CONTROLLED VOCABULARIES WITH OWL ONTOLOGY: A CASE STUDY FROM THE DOMAIN OF MOLECULAR INTERACTIONS*

MELISSA J. DAVIS

*Institute for Molecular Bioscience and ARC Center of Excellence in Bioinformatics,
University of Queensland, St. Lucia
Brisbane, QLD 4076, Australia*

ANDREW NEWMAN

*School of Information Technology and Electrical Engineering,
University of Queensland, St. Lucia
Brisbane, QLD 4076, Australia*

IMRAN KHAN

*School of Information Technology and Electrical Engineering,
University of Queensland, St. Lucia
Brisbane, QLD 4076, Australia*

JANE HUNTER

*School of Information Technology and Electrical Engineering,
University of Queensland, St. Lucia
Brisbane, QLD 4076, Australia*

MARK A. RAGAN

*Institute for Molecular Bioscience and ARC Center of Excellence in Bioinformatics,
University of Queensland, St. Lucia
Brisbane, QLD 4076, Australia*

Many efforts at standardising terminologies within the biological domain have resulted in the construction of hierarchical controlled vocabularies that capture domain knowledge. Vocabularies, such as the PSI-MI vocabulary, capture both deep and extensive domain knowledge, in the OBO (Open Biomedical Ontologies) format. However hierarchical vocabularies, such as PSI-MI which are represented in OBO, only represent simple parent-child relationships between terms. By contrast, ontologies constructed using the Web Ontology Language (OWL), such as BioPax, define many richer types of relationships between terms. OWL provides a semantically rich structured language for describing classes and sub-classes of entities and properties, relationships between them and expressing domain-specific rules or axioms that can be applied to extract new information through semantic inference. In order to fully exploit the domain knowledge inherent in domain-specific controlled vocabularies, they need to be represented as OWL-DL ontologies, rather than in formats such as OBO. In this paper, we describe a method for converting OBO vocabularies into OWL and

* This work has been performed as part of a collaborative research and development project with Pfizer Global R&D and is partially supported by ARC grand CEO348221.

class instances represented as OWL-RDF triples. This approach preserves the hierarchical arrangement of the domain knowledge whilst also making the underlying parent-child relationships available to inference engines. This approach also has clear advantages over existing methods which incorporate terms from external controlled vocabularies as literals stripped of the context associated with their place in the hierarchy. By preserving this context, we enable machine inference over the ordered domain knowledge captured in OBO controlled vocabularies.

1 Introduction

Molecular biology as a field encompasses several dynamic sub-domains undergoing rapid expansion with resultant rapid discovery and growth in acquired data. High-throughput techniques and large scale biological research, such as genome and transcriptome sequencing projects and expression studies, generate abundant data. However, a significant gap exists between data acquisition and knowledge discovery. These massive quantities of data are frequently produced through a distributed effort, and need to be integrated for analysis and final presentation [1, 2]. Likewise, techniques such as expression profiling produce large quantities of raw data which must be recorded and described [3]. In addition to data exchange and integration issues, many projects in computational and systems biology focus on the analysis of high-level properties of biological systems. Such projects might, for example, compare the distribution of protein functional classes between genomes [4], or analyse the genetic regulatory network of an organism [5]. Such analysis requires the integration of heterogeneous information produced from multiple sources at varying levels of resolution and described using highly variable terminologies. Solutions to these exchange and integration challenges include provision of the data in delimited text files, databases, and XML documents conformant with a given XML schema. The semantic meaning of the data is not however explicit within the documents, and relies on some external definition of concepts and relationships in the data.

Analysis at the level of biological systems also requires reasoning over large and complex data sets is beyond the ability of humans. Machine reasoning has the ability to uncover implicit relationships in the data, rather than simply retrieving explicitly represented data, as is the case of querying a database. However, machine reasoning over large and complex data sets requires the use of appropriate and meaningful knowledge representations of the domain area combined with presentation of the data in machine-readable format [6]. One technique for implementing knowledge representation that has been readily adopted in biology has been the construction of bio-ontologies to establish a precisely (if not formally) defined way to model and express the knowledge of a domain in terms of defined concepts: the classes of “things”, the relationships that exist between these classes, and the rules or axioms that apply to these concepts in the domain. The need for ontology development was recognized almost a decade ago in the creation of perhaps the most widely adopted biological ontology, the Gene Ontology (GO) [7].

An important consideration is why choose to create or use an ontology over traditional database schema, which are widely adopted knowledge representations in the

field of molecular biology [8]. While there have been few successful demonstrations of automated inference using bio-ontologies, there are significant reasons for their adoption [9]. The key reason is that ontologies are designed to evolve over time and to facilitate integration of data, while database schemas are not [10]. Database schemas are typically considered an internal design decision for a given application and rarely, if ever, are schemas from other databases reused. A specific ontology, on the other hand, is an external, global resource that is meant to be reused, extended and integrated with other ontologies. An ontology is also more expressive than a database [11]. Finally, databases rarely allow the preservation of data. It is still common simply to add attributes to an existing schema rather than splitting it logically due of the extent of the data migration [12]. Ontologies provide a separation between the actual data and the metadata or descriptions of the datasets and their relationships. This allows the data to be migrated independently from changes within the ontology [10].

Many types of knowledge representation exist, and there are many views of what constitutes machine reasoning [13]. The majority of ontologies developed in the biological domain to date do not take advantage of this background [14]. Technologies developed by the World Wide Web Consortium (W3C; <http://www.w3.org/>) to support machine reasoning include the Resource Description Framework (RDF) and Web Ontology Language (OWL). Both RDF and OWL support machine inference across resources on the web.

While some bio-ontologies have been constructed using these standards, or have been converted into a form compliant with these standards [15], most do not take advantage of the W3C recommendations [14] (see OBO Foundry <http://www.obofoundry.org>). Many are presented as controlled vocabularies where concepts are represented taxonomically, and relationships are predominantly “is-a” or “part-of” relationships that establish the tree-like structure of the vocabulary. While these structures create well ordered catalogues of concepts relevant to a domain, they do not typically allow for the expression of rules defining other relationships between concepts. The end result is a simplified, flattened model of the domain that lacks the semantic depth or logical support to enable a reasoner to infer new relationships or new information.

A significant challenge in molecular biology is to understand the molecular interactions that occur within cells; research in cell and structural biology shows that many proteins rely on a complex network of interacting partners to achieve their correct localization and functional state in the cell. High-quality molecular interaction data are largely described in journal articles using natural language. Because of the unstructured nature of the observations, the discipline of molecular interactions is covered by several overlapping ontologies. Some of these, such as BioPax (<http://www.biopax.org>) and the Protein Standards Initiative Molecular Interaction vocabulary (PSI-MI; <http://www.psidev.info/>) provide significant coverage over concepts relevant to the domain, while others, such as GO, and the NCBI-Taxonomy (<http://www.ncbi.nlm.nih.gov/>) intersect with the field. The diverse formats of these

overlapping ontologies make molecular interactions a useful test domain for strategies to integrate bio-ontologies and reuse domain knowledge.

2 Results

The field of bio-ontology development is active, and already populated with a number of general and domain specific ontologies that have been developed, or are under development. We reviewed ontologies listed by the OBO Foundary; and the National Center for Biomedical Ontology (NCBO; <http://www.bioontology.org/>). Of the ~70 ontologies listed at these sites, around three quarters are written using the OBO format [16], with the remainder using other formats, including OWL, Protégé files and plain text.

OWL is specifically designed to construct ontologies that support machine reasoning [11]. For this reason, we choose to use OWL DL (description logics) to construct a high-level ontology to integrate concepts from relevant biological ontologies and vocabularies not expressed in OWL. Given that knowledge acquisition is one of the most time consuming and necessarily manual parts of ontology construction, the knowledge captured in non-OWL ontologies constitutes a valuable resource. One approach suggested in such cases is to construct a new ontology using OWL [14], however this under estimates the value of knowledge represented in other formats. Practical strategies to rescue domain knowledge captured in non-OWL ontologies will have obvious applicability in a domain such as molecular biology where the majority of vocabularies are not expressed in OWL.

We have reviewed two ontologies used to describe molecular interactions: the OWL ontology BioPax, and the OBO ontology PSI-MI. However it is not our intention to comparatively evaluate these ontologies, as has been done recently [17, 18]. Briefly, BioPax is designed to describe pathway rather than specific molecular interaction data. However of the ~40 classes and ~70 properties that BioPax defines, many are key concepts and relationships necessary to describe molecular interactions. The PSI-MI vocabulary, on the other hand, is specifically designed to describe molecular interaction data and captures >800 concepts from the domain. However it is represented in OBO and expresses only hierarchical relationships between these classes. While BioPax lacks the descriptive power of PSI-MI, it is more suited for machine reasoning because it is represented in OWL. BioPax recognizes the value of external controlled vocabularies such as PSI-MI and GO, by providing a facility to exploit these external vocabularies through the inclusion of a class `openControlledVocabulary`. This class stores a term from an external vocabulary along with a cross reference holding the identifier of that term and the name of the vocabulary (as literal strings).

However, it is not sufficient merely to store the data from external controlled vocabularies as literals - the term becomes devoid of meaning if taken out of the context of the original concept tree. In order to preserve the meaning of the term its relationship to other terms throughout the hierarchy, must also be preserved. For example, most

biologists would know that “mouse” is not only a label for an organism, but also that an organism which is a *mouse* is also a *mammal* and a *vertebrate*. If “mouse” is used as a text-string label, it lacks the meaning a biologist would ascribe to the concept, and all that is available is the string of text. While applications could be written that interpret the string “mouse” to have a particular meaning, there is nothing in the representation that makes the meaning explicit. However, if an OWL ontology containing a class hierarchy of taxonomic terms was used in the representation, then an organism that was an instance of the class *mouse* would inherit class membership of super-classes in the hierarchy as instances of each sub-class are also instances of super-classes. This allows an instance of *mouse* to be recognized as an instance of *mammal* and *vertebrate*. There is no need for any special encoding of this information into applications interpreting the data, since the meaning is now made explicit in the ontology itself. The meaning of the term is preserved through maintaining its relationship to other terms and its place in the hierarchy.

To illustrate this point, consider the natural language expression, “The protein Emerin is localized to the nuclear inner membrane” (Figure 1).

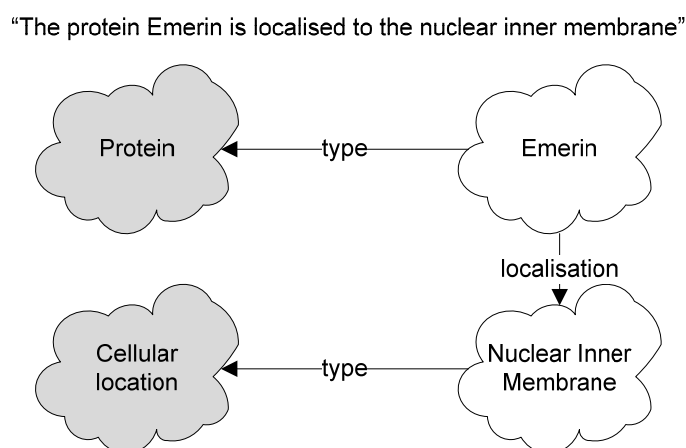


Figure 1. A natural language assertion decomposed into generic concepts (shaded) and specific instances of those concepts. The specific instance “Emerin” is of the generic type *protein*, while the specific instance “Nuclear inner membrane” is of the generic type *cellular location*. The relationship between “Emerin” and its localisation is represented by a labeled arrow, *localisation*.

This statement is composed of generic and specific concepts and relationships: two generic concepts, protein and cellular location, and two specific instances of these concepts, “Emerin”, a protein, and “nuclear inner membrane”, a cellular location. A relationship also exists between these two instances, namely, that “Emerin” has a property localisation, the value of which is “nuclear inner membrane”. The same statement could also be expressed using the BioPax ontology, as illustrated in Figure 2, by using the openControlledVocabulary class to include a term from an external vocabulary like the GO Cellular Component hierarchy.

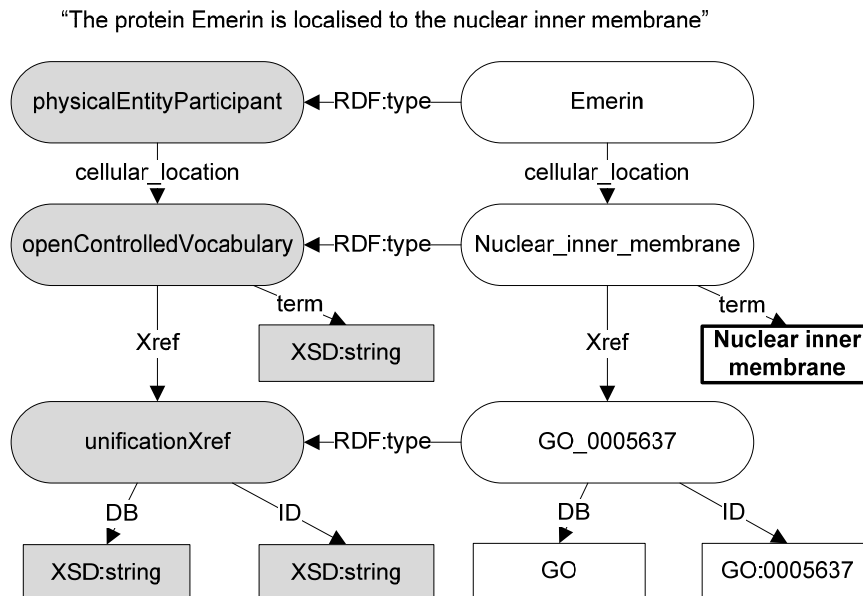


Figure 2. Expression of an assertion of subcellular location in BioPax. Classes and properties from the BioPax ontology are shaded on the left of the diagram, while the instance data describing the localisation of Emerin is on the right. Elipses represent classes, or instances of classes, while rectangles represent typed (string) literals. The string in bold, “nuclear inner membrane” has been imported from the GO Cellular Component hierarchy.

However, as the references to the term from the external vocabulary are all simple text strings, they lack context or meaning. A search for proteins where the value of `cellular_location` is the string “organelle membrane” would not retrieve proteins where this value was “nuclear inner membrane” unless the query application was hard-coded with additional information about the relationship between these two strings. External terms used in this fashion lack meaning. Because of the limitation of the BioPax approach, we developed a different approach to include domain knowledge captured in external controlled vocabularies (see Figure 3).

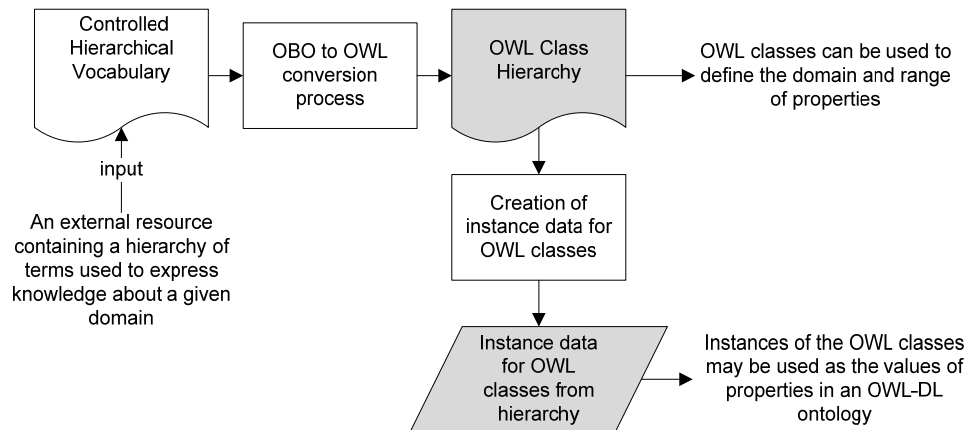


Figure 3. Integration process for an OBO Controlled Vocabulary. The OBO vocabulary is converted into OWL-DL. This results in an OWL class hierarchy where terms from the original vocabulary become OWL classes. Instance data is then created for the ontology so each class has a single instance comprising the vocabulary term that can then be used as an object or subject of triples. Outputs from this process are shaded.

An external controlled vocabulary that contains relevant descriptive terms is converted into OWL-DL. Most frequently, the external vocabulary is in OBO format, so we currently use one of several OBO to OWL conversion applications [19-21]. However this process may be more generally applied to any hierarchical controlled vocabulary such as NCBI Taxonomy, which is not in OBO format. It is important that both the class hierarchy and the instance data for this hierarchy are created. OWL DL classes are used to define restrictions for properties within the ontology, through the specification of allowable domain and range values [11]. However, a class cannot also be an instance, and only instances may be used as the values of properties. For this reason, a single instance of each class is created, taking the form of the original term from the hierarchical vocabulary. At the end of this process, both an OWL ontology representing the terms from the controlled vocabulary and a set of instance data are available (see Figure 3) to use in conjunction with an OWL ontology, as illustrated in Figure 4.

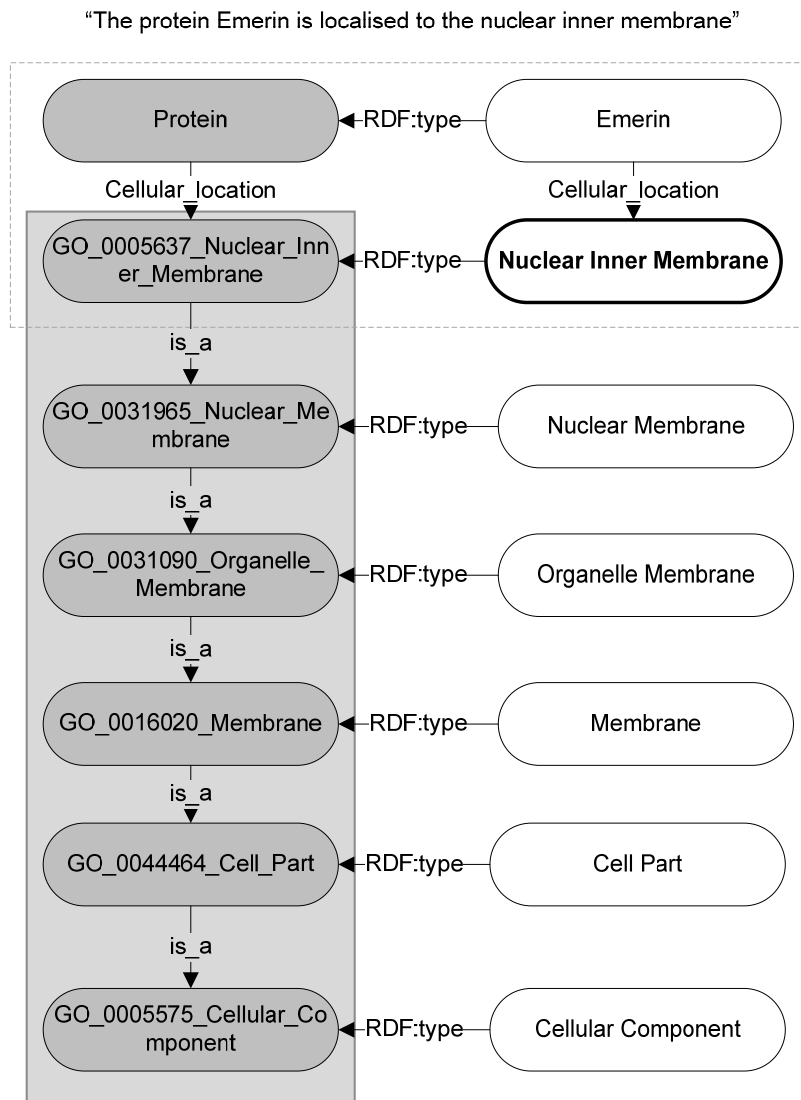


Figure 4. Application of the conversion process. Classes and instances used to express the statement are boxed in a dashed line. The term, “nuclear inner membrane” in bold is an instance of the class `GO_0005637_Nuclear_Inner_Membrane`, which is related to other classes in the hierarchy shaded box. Domain knowledge is explicitly captured in these hierarchical relationships, so that the relationship of the protein “Emerin” to other cellular locations can be inferred.

In this example, the instance “nuclear inner membrane” is used as the value of the `cellular_location` property. Not only is this value meaningful to a person who understands what is meant by the words, it is also meaningful to a machine reasoner that has access to the underlying ontology. This machine reasoner, when presented with the assertion that

Emerin is located in the “nuclear inner membrane”, may correctly infer that Emerin is also located in an “organelle membrane”, and located in a “membrane”.

By using this process to incorporate components of external vocabularies under a high-level extensible ontology, external terms become more than text labels, and enable implicit relationships to be extracted from explicit data.

3 Discussion

Many of the existing bio-ontologies are written in the OBO format, and represent a rich source of biomedical domain knowledge. By using the approach described here, vocabularies covering specific aspects of a domain may be plugged into a high level extensible OWL ontology designed to facilitate these modular extensions. The parent-child relationships of the new vocabulary are maintained, and made available to a reasoner to infer implicit relationships from the explicitly represented data. The current strategy used by BioPax is inadequate for machine reasoning. To make bio-ontologies useful for machine reasoning, they need to be explicitly represented in languages such as OWL. Converting and importing relevant hierarchies of terms and associated instances is one solution to maintaining the meaning of terms in controlled vocabularies.

We have constructed a high-level ontology that imports classes and properties from BioPax to describe such things as entities, references and cross-references. To avoid using the BioPax `openControlledVocabulary` class to incorporate external information, we construct new properties with BioPax classes for domains, and ranges that specify classes in the converted OWL class hierarchy targeted by the property. For example, we create a property `Cellular_location`, which has a range of `GO_0005575_Cellular_Component`. While the property conceptually maps to the BioPax property of the same name, the range is restricted to elements from the converted GO Cellular Component hierarchy, that is, all the subclasses of the class `GO_0005575_Cellular_Component`. In the same way, we can exploit the rich vocabulary in PSI-MI that describes experimental methods used to determine molecular interactions by creating a property `Experimental_method` (which conceptually maps to the BioPax property `experimental_form`) which has a range restricted to values taken from subclasses of the converted PSI-MI `MI_0045_experimental_interaction_detection` class describing experimental methods. This strategy enables us to extend the ontology to incorporate descriptive vocabulary terms from external controlled vocabularies, without losing the context, or meaning, that those terms have in their vocabulary of origin.

One concern when creating an ontology or expanding an existing ontology is the trade off between the ability of the ontology to express concepts in the domain, and to provide tractable inference. The effects of this trade off are difficult to evaluate [23]. A strategy which we will explore to address this is to identify which branches of a given concept tree are required and include only those branches.

Since so many biological concepts are framed in terms of hierarchically inherited properties, and the majority of biological ontologies take the form of hierarchical

controlled vocabularies, the process described here is a useful generic strategy for incorporating the existing wealth of ordered knowledge into a semantically rich ontology constructed using OWL. This will help to extend the utility of bio-ontologies into the arena of machine inference.

Acknowledgments

The authors thank Pfizer Inc. for support of the BioMANTA project. We also thank our colleagues in the BioMANTA project for useful suggestions and discussion.

References

1. F. Collins, M. Morgan and A. Patrinos, *Science*, **300**, 286 (2003).
2. P. Carninci et al. *Science*, **309**, 1559 (2005).
3. A. Brazma et al. *Nat. Genet.* **29**, 365 (2001).
4. M. A. Andrade et al., *J. Mol. Evol.* **49**, 551 (1999).
5. S. Li et al., *Science*, **303**, 540 (2004).
6. F. Bry and M. Marchiori, in *Proc. 2nd EWIMT*, London, UK, IEEE. (2005).
7. M. Ashburner et al., *Nat. Genet.* **25**, 25 (2000).
8. M. Y. Galperin, *Nucleic Acids Res.* **35**, D3 (2007).
9. M. Keet, M. Roos and M. Marshall, in *Proc. 3rd OWLED*, Innsbruck, Austria, CEUR-WS, (2007).
10. N. Noy and M. Klein, *KAIS*, **6**, 428 (2004).
11. S. Bechhofer et al., in *W3C Recommendations*, (2004).
12. R. Elmasri, S. Navathe and C. Shanklin, *Fundamentals of Database Systems*, Addison-Wesley, Boston (2000).
13. R. Davis, H. Shrobe and P. Szolovits, *AI Magazine*, **14**, 17 (1993).
14. L. Soldatova and R. King, *Nat. Biotech.* **23**, 1095 (2005).
15. M. E. Aranguren et al., *BMC Bioinformatics*, **8**, 57 (2007).
16. R. G. Cote et al., *BMC Bioinformatics*, **7**, 97 (2006).
17. L. Stromback and P. Lambrix, *Bioinformatics*, **21**, 4401 (2005).
18. L. Stromback et al., *Brief. Bioinformatics*, **7**, 331 (2006).
19. D. A. Moreira and M. A. Musen, *Bioinformatics*, **23**, 1868 (2007).
20. C. Goldbreich and I. Horrocks, in *Proc. 3rd OWLED*, Innsbruck, Austria, CEUR-WS, (2007).
21. S. H. Tirmizi and D. P. Miranker, *Technical Report*
<http://www.cs.utexas.edu/~hamid/pub/tirmizi-obo2owl-tr-06-47.pdf>
22. D. Martin et al., *Genome Biol.* **5**, R101 (2004).
23. H. Levesque and R. Brachman, *Comput. Intell.* **3**, 78 (1987).