Towards an Ontology for Data-driven Discovery of New Materials

Kwok Cheung[1], John Drennan[1], Jane Hunter[2]
[1]AIBN, The University of Queensland
St Lucia, Queensland, Australia
[2]ITEE, The University of Queensland
St. Lucia, Queensland, Australia
(kwokc, jane)@itee.uq.edu.au, j.drennan@uq.edu.au

Abstract

Materials scientists and nano-technologists are struggling with the challenge of managing the large volumes of multivariate, multidimensional and mixed-media data sets being generated from the experimental, characterisation, testing and post-processing steps associated with their search for new materials. In addition, they need to access large publicly available databases containing: crystallographic structure data; thermodynamic data; phase stability data and ionic conduction data. Materials scientists are demanding data integration tools to enable them to search across these disparate databases and to correlate their experimental data with the public databases, in order to identify new fertile areas for searching. Systematic data integration and analysis tools are required to generate targeted experimental programs that reduce duplication of costly compound preparation, testing and characterisation. This paper presents MatOnto – an extensible ontology, based on the DOLCE upper ontology, that aims to represent structured knowledge about materials, their structure and properties and the processing steps involved in their composition and engineering. The primary aim of MatOnto is to provide a common, extensible model for the exchange, re-use and integration of materials science data and experimentation.

Introduction and Objectives

The advent of high-throughput, combinatorial and robotic laboratory instruments, atomic resolution microscopes and high speed modelling and simulation software tools is triggering an explosive growth in the magnitude and complexity of materials data. Materials science data ranges from complex compound preparation and processing workflows, to spectrographic analyses, 2D nano-scale microscopy images, textual publications, numerical data, animations, 3D crystallographic structures and complex phase diagrams.

Materials informatics is emerging as a new discipline addressing the issues of data management, curation, integration and analysis that are challenging materials scientists. Materials informatics is defined as *the high speed robust acquisition, management, analysis and dissemination of diverse materials data*. Materials data access, acquisition, interoperability and curation were recently identified as critical cyberinfrastructure imperatives for the materials science community [1, 2].

Critical requirements include: persistent unique identifiers for materials science resources; metadata standards for describing samples, processes, properties; common semantic models/ontologies to enable mapping between database schemas, information integration and semantic interoperability; laboratory information and provenance capture systems that capture the processes both in the laboratory as well as in the post-processing of the data. Semantic Web technologies are essential to addressing these issues, and the development of Materials Ontology (MatOnto) is a significant step towards an integrated solution.

Ontologies provide rich machine-processable semantic descriptions; formal definitions of domains by defining classes, properties and relationships between them in Web Ontology Language (OWL); and a basis to enable reasoning and deduction of new information. Ontologies enable semantic interoperability between resources, services, databases, and devices via inter-related knowledge structures.

The remainder of this paper describes the MatOnto ontology, which aims to:

- Provide an extensible framework that encapsulates the top level structured knowledge of materials science;
- Enable integration of and mapping between disparate databases within the materials science domain;
- Enable the modelling and capture of precise provenance data in both the digital and physical domains. This is essential to enable verification, validation, comparison and re-use of experimental results;
- Enable the inferencing and extraction of new knowledge in the materials science domain, through the application of SWRL rules and a reasoning engine.

Example Scenario

At the University of Qld, we are working with fuel cell scientists who are searching for novel oxygen ion conducting materials that can operate more efficiently at lower temperatures for longer durations. The electrolyte

compound must have oxygen conductivities > 10^{-1} Scm^{-1} and mechanical and chemical stability at elevated temperatures (500 °C). Based on past experience and intuitive knowledge, fuel cell scientists want rapid answers to queries such as: "Give me compounds that contain tungsten-oxygen-X (where X is a different cation), with bond lengths between Y and Z nm, with large anomalies and anisotropy in the positional parameters of oxygen, with bond angles between J° and K° and which are stable below 500 °C. Figure 1 below illustrates the scientific workflow and methodology for this research project.

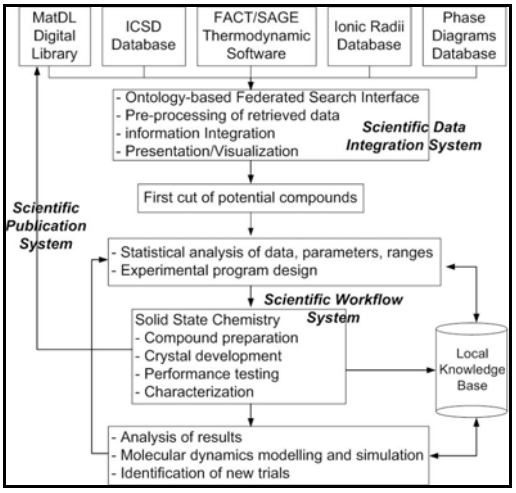

Figure 1: Schematic illustration of overall project structure and components

To answer such queries, scientists currently have to manually search, retrieve, process and correlate data from a number of related but disparate databases including: the Inorganic Crystal Database (ICSD) [3]; the thermodynamic FactSage database [4]; the Ionic Radii Database [5]; and NIST Phase Equilibria Diagrams database [6]. One of the greatest hurdles to this process is that the search interfaces, metadata terms, data structures, formats and metrics are inconsistent across these databases. For example, *temperature factors* can be represented in three different formats (Isotropic Temperature Factor (ITF), temperature factor (β) and the mean square amplitude of vibration (U)). Sophisticated database integration and mining tools are required so fuel cell scientists can more easily retrieve answers to such queries. Fast and intelligent tools are required to seamlessly interrogate, retrieve, integrate and present data so users can iteratively hone in on areas of interest. XML schemas (e.g., the Materials Markup Language MatML [7]) can specify the data structures and formats within each database. However, machine-processable ontologies are required to provide the semantic mappings between related terms and to dynamically pre-process, integrate, correlate and reason across data retrieved from the disparate databases.

Related Work

There have been a number of independent efforts in ontology development within the materials science community. These prior efforts have focussed on: 1) developing ontologies for data integration; 2) extracting knowledge from text; 3) a cross-disciplinary classification scheme for nanoscale research; and 4) demonstrating ontology development based on object-oriented database schemas.

Ashino et al [8] presents an ontology for the material selection process, that aims to integrate material properties such as the creep property, with materials databases via a standardized XML schema. Second, The PLINIUS ontology [9] was designed specifically for the domain of ceramic materials and was aimed at semi-automatic knowledge extraction from texts. Tanaka [10] presents a meta-level ontology as a multi-disciplinary classification scheme within Nanoscience and Nanotechnology. Finally, Ono et al. [11] argues that a set of object-oriented classes in a domain provides an ideal framework for ontology development due to the high similarity between the specifications of an ontology and an object-oriented system.

Ashino's and PLINIUS's ontologies focus on the details of specific sub-disciplines within materials science while Tanaka provides a coarse classification of sub-disciplines within nanoscale research. All of them fail to mention the importance of extensibility of their ontologies. In contrast, MatOnto has been developed as a high-level materials science ontology in order to address this issue. As a result of this design, MatOnto provides a platform for integration of existing or emerging sub-disciplinary ontologies within materials science domain. It also provides a platform for linking to relevant cross-disciplinary data e.g., bio-materials.

Ashino's and Tanaka's ontologies are primarily extracted from materials data standards or prestigious data sources. While Ono's proposed phase diagram ontology includes reinvented terms for new concepts. The PLINIUS ontology develops complex concepts based on atomic concepts and construction rules. The PLINIUS approach may work within a limited scope of a problem domain, but it is doubtful that such an approach will satisfactorily capture every term within the materials science domain accurately. Consequently, we chose to develop MatOnto by merging Ashino's and Tanaka's approaches with DOLCE – and also calling on existing XML standards such as MatML [7].

MatOnto Development

MatOnto's design principles are to provide an ontology that: 1) is based on an upper ontology, an advanced knowledge representation system, that is *a library of richly structured and well-understood abstract data types and structural organizational principles, which make the technical aspects of ontology construction easier and more reliable* [12]*;* 2) leverages existing peer-reviewed ontologies or vocabularies developed through community

consensus; 3) enables integration of those high priority databases identified by our fuel cell collaborators. Below we describe the six steps in the process of developing the MatOnto ontology.

Firstly, we decided to use DOLCE [13] the upper ontology developed by the Laboratory for Applied Ontology (LOA), as the upper basis for MatOnto. DOLCE stems from the *Entity* root class. *Entity* has three subclasses: E*ndurant, Perdurant* and *Abstract*, from which we define MatOnto subclasses.

Secondly we leveraged a number of existing peer-reviewed ontologies and a classification system: Ontolingua's Standard Units and Dimensions [14]; the Joint Academic Classification of Subjects (JACS) [15] ; W3C's Time Ontology in OWL [16]; and AIFB's Semantic Web for Research Communities (SWRC) ontology [17] .

Thirdly, we extended EXPO [18] an ontology for describing scientific experiments with the ABC Metadata Ontology [19] in order to enrich EXPO with the concepts of *events* and *processes*. EXPO is primarily a taxonomy of scientific experiments, while the ABC Ontology models events in both the physical domain and a digital object's lifecycle.

Fourthly, we developed the top-level ontology for materials science according to [20] and [21], beginning with class *matonto:Material*, which is linked to *jacs:Materials Science* of the Joint Academic Classification of Subjects.

Figures 2 and 3 represent the complete top-level view of MatOnto - the classes with prefix *dolce* are from DOLCE. Figure 2 also shows the use of classes (with prefixes *ontolingua, swrc* and *w3c* respectively) from existing peer-reviewed ontologies [14,16,17]. Figure 3 illustrates the use of classes from JACS (those with prefix *jacs*) and EXPO [18] (prefix *expo*), and the root class of our MatOnto ontology – *matonto:Material*.

Figure 4 demonstrates the merging of classes from the EXPO and ABC ontology. While Figure 5 demonstrates the properties and relationships of the core class *expo:Scientific Activity which* has been extended using MatOnto classes to model materials science experiments.

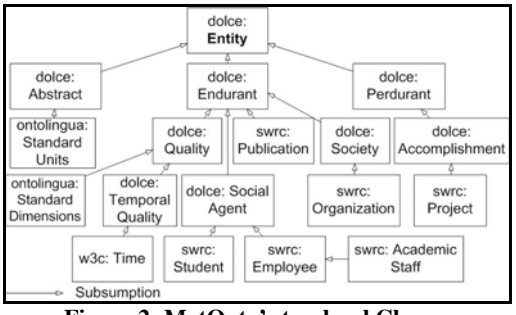

Figure 2: MatOnto's top-level Classes

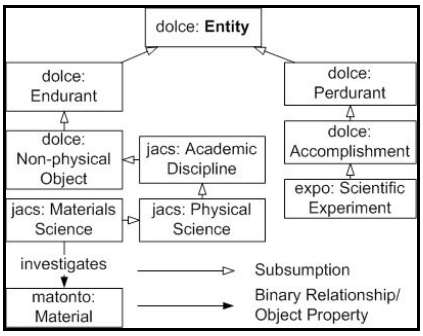

Figure 3: Materials Science Extensions to the Top-level Classes

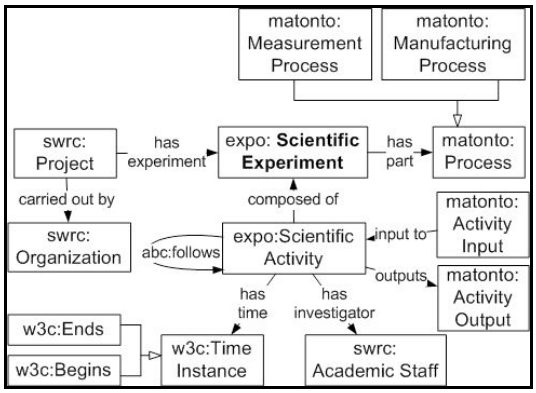

Figure 4: The merged EXPO and ABC ontologies

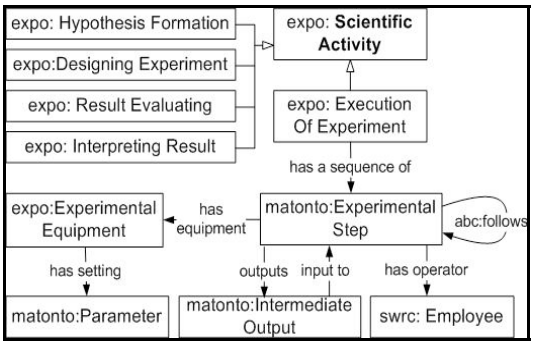

Figure 5: The view stemming from the *Scientific Activity* Class

We identified five core properties associated with *matonto:Material*:

1) *matonto*:*Property* – the materials properties shown in Figure 6;
2) *matonto:Family* – the materials classification shown in Figure 7;
3) *matonto*:*Process* – the materials manufacturing and measurement processes shown in Figure 8;
4) *matonto*:*Structure* – the materials' structure, also shown in Figure 8; and
5) *matonto:Measurement* – the data resulting from measurement or characterisation process shown in Figure 9. We have drilled down to certain levels and structured the associated concepts in a logical way.

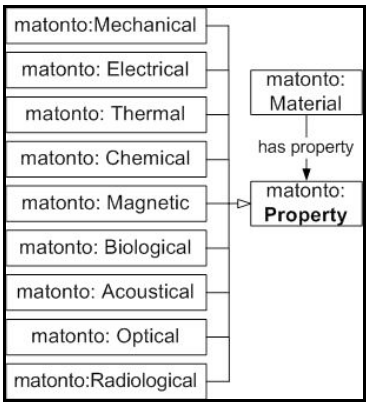

Figure 6: Materials Property

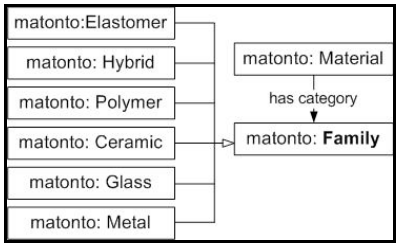

Figure 7: Materials Family

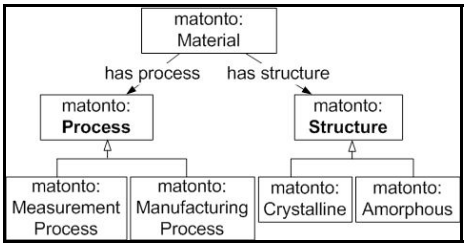

Figure 8: Materials Processes and Structures

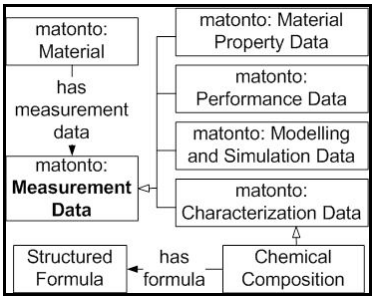

Figure 9 Materials Measurement Data

Fifthly, we developed a sub-disciplinary ontology describing the concepts associated with crystalline structures according to Crystallographic Information Framework [22] and subsumed it under the class *matonto:Crystalline*, which is a sub-class of *matonto:Structure*.

Figures 10 and 11 demonstrate the complete top-level view of the Crystalline Structure Ontology.

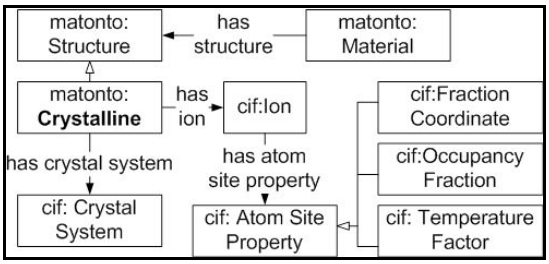

Figure 10: Crystalline Structure Ontology - ionic data and crystal systems

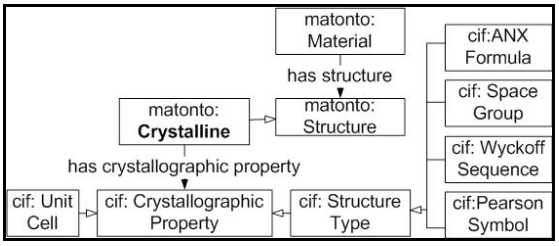

Figure 11: Crystalline Structure Ontology - crystallographic properties

Finally we developed a Scientific-Data ontology to describe the different types of numerical and multimedia data. Figures 12 and 13 illustrate the complete high-level view of the MatOnto Scientific-Data ontology.

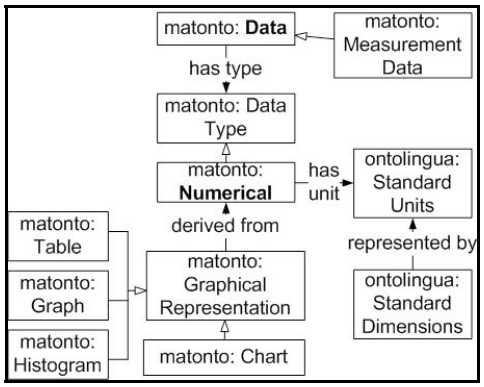

Figure 12: Scientific-Data Ontology - numerical data

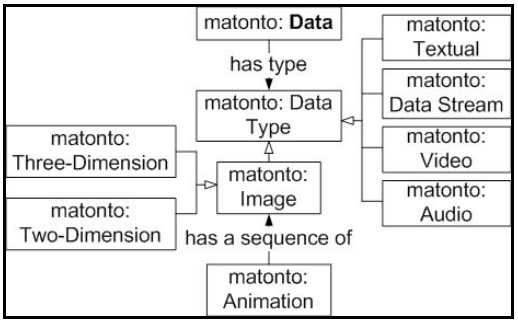

Figure 13: Scientific-Data Ontology - multi-media data

Linking of all of these sub-ontologies via their common classes generates the complete MatOnto ontology.

Discussion

Achievements

MatOnto satisfies the objectives outlined in the *Introduction and Objectives* section. It enables integration of existing related sub-disciplinary and relevant ontologies through the top level materials science classes shown in Figures 6 to 9. The Crystalline Structure Ontology shown in Figures 10 and 11 enables the integration of and mapping between, the Inorganic Crystal Structure Database (ICSD) [3] and Ionic Radii databases [5]. The extended EXPO Ontology shown in Figures 4 and 5 enables the capture of precise provenance data and the inferencing of new knowledge (e.g., relationships between nodes that are not explicitly related). This aspect is used to automatically infer coarse-grained views of the scientific methodology from fine grained provenance trails, for publication or elearning purposes [25].

Ontological Assessment

MatOnto's quality has been assessed based on Gruber's five criteria [23]: clarity, coherence, extendibility, minimal encoding bias, and minimal ontological commitment - with satisfactory results. First, MatOnto possesses clarity because its vocabulary is sourced from peer-reviewed ontologies and existing standardized taxonomies. Secondly, MatOnto does not have incoherency issues because no concepts are derived via inferencing. Thirdly, MatOnto is extensible because DOLCE together with JACS provide a proven platform for integrating disciplinary ontologies. The high-level materials science ontology provides a platform for integrating sub-disciplinary ontologies within materials science domain, e.g. the Crystalline Structure Ontology. Fourthly MatOnto has no encoding bias because it is free of implementation details. Finally MatOnto has low ontological commitment because we have reused existing peer-reviewed ontologies and extended them based on standardized vocabularies.

Evaluation

We have represented MatOnto in the Web Ontology Language (OWL) and we are in the process of evaluating it through its application within three software tools that we have developed and that are being user tested by our materials science collaborators: 1) a federated ontology-based search interface that enables materials scientists to search, retrieve and integrate data from the ICSD, Ionic Radii and Phase Diagrams database; 2) a scientific workflow system [24] that collects scientific results with provenance data during a fuel-cell manufacturing process, 3) SCOPE [25] a Scientific Compound Object Publishing and Editing tool that generates OAI-ORE [26] compliant compound objects. This tool enables the visualization and exploration of provenance trails by expanding or collapsing links between nodes in the scientific workflow. A set of SWRL rules are being developed, specifically for materials science, that can be executed using the Pellet reasoning engine [27], to infer new implicit relationships and knowledge from explicit data. Inferencing is applicable to a number of aspects of materials science, including:

- Inferring relationships between processing parameters and structure
- Inferring relationships between structure and properties or behaviour
- Inferring structural features from automatic image analysis of microscopy images
- Inferring coarse grained views of provenance from fine-grained provenance trails.

We plan to further explore the application of semantic inferencing to knowledge extraction from materials science data, in the near future.

Conclusions

In this paper, we have described MatOnto - an ontological framework that encapsulates the knowledge structure of materials science and that can be easily extended to integrate with related ontologies. MatOnto enables materials scientists to search, retrieve and integrate data from heterogeneous and disparate data sources, based on a common set of ontological terms. It also enables the capture of processing steps and provenance information both within the laboratory as well as within the computing environment. This enables the repeatability, exchange, comparison and re-use of experimental results. The MatOnto ontology also provides the potential for inferencing and extraction of new knowledge using SWRL rules defined by domain experts (e.g., fuel cell scientists) and a reasoning engine (e.g., Pellet). MatOnto provides an essential and fundamental component of the cyberinfrastructure requirements of the materials science community.

References

1. Billinge, S.J.L., K. Rajan, and S.B. Sinnott. From Cyberinfrastructure to Cyberdiscovery in Materials Science: Enhancing outcomes in materials research, education and outreach. in Report from NSF-sponsored workshop held in Arlington, Virginia. 2006.

2. Hunt, W., Materials informatics: Growing from the Bio World. JOM Journal of the Minerals, Metals and Materials Society, 2006. 58(7): p. 88-88.

3. Belsky, A., et al., New developments in the Inorganic Crystal Structure Database (ICSD): accessibility in support of materials research and design. Acta Crystallographica Section B, 2002. 58(3 Part 1): p. 364-369.

4. Bale, C.W., et al., FactSage thermochemical software and databases. Calphad, 2002. 26(2): p. 189-228.

5. Shannon, R., Revised effective ionic radii and systematic studies of interatomic distances in halides and chalcogenides. Acta Crystallographica Section A, 1976. 32(5): p. 751-767.

6. NIST Phase Equilibria Diagrams Database. [cited 14 Jan 2008]; Available from: http://www.ceramics.org/publications/phase.aspx.

7. MatML a new markup language for materials information.(ASM NEWS). Advanced Materials & Processes, 2005. 163(5): p. 62(1).

8. Ashino, T. and M. Fujita, Definition of a web ontology for design-oriented material selection. Data Science Journal, 2006. 5: p. 52-63.

9. van der Vet, P.E., P.-H. Speel, and N.J.I. Mars. The PLINIUS Ontology of Ceramic Materials. in the Eleventh European Conference on Artificial Intelligence (ECAI'94) Workshop on Comparison of Implemented Ontologies. 1994. Amsterdam: The Netherlands.

10. Tanaka, M. Toward a Proposed Ontology for Nanoscience. in CAIS/ACSI 2005: Data, Information, and Knowledge in a Networked World. 2005. The University of Western Ontario, London, Ontario.

11. Ono, N., et al. Ontology for phase diagram databases. in Intelligent Processing and Manufacturing of Materials, 1999. IPMM '99. Proceedings of the Second International Conference on. 1999.

12. Colomb, R.M., Formal versus Material Ontologies for Information Systems Interoperation in the Semantic Web. The Computer Journal, 2006. 49(1): p. 4-19.

13. Aldo, G., et al., Sweetening Ontologies with DOLCE, in Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management. Ontologies and the Semantic Web. 2002, Springer-Verlag.

14. Gruber, T.R. and G.R. Olsen. An Ontology for Engineering Mathematics. in Fourth International Conference on Principles of Knowledge Representation and Reasoning. 1994. Gustav Stresemann Institut, Bonn, Germany: Morgan Kaufmann.

15. The Joint Academic Classification of Subjects. 2000 [cited 1 November 2007]; Available from: http://www.hesa.ac.uk/dox/jacs/JACS_complete.pdf.

16. Time Ontology in OWL - W3C Editor's Draft. 2005 20 September 2005 [cited 1 November 2007]; Available from: http://www.isi.edu/~pan/SWBP/time-ontology-note/time-ontology-note.html.

17. Sure, Y., et al., The Ontology – Semantic Web for Research Communities, in Progress in Artificial Intelligence. 2005. p. 218-231.

18. Soldatova, L.N. and R.D. King, An ontology of scientific experiments. Journal of The Royal Society Interface, 2006. 3(11): p. 795-803.

19. Lagoze, C. and J. Hunter, The ABC Ontology and Model. Journal of Digital Information, 2001. 2(2).

20. Ashby, M., H. Shercliff, and D. Cebon, Materials : engineering, science, processing and design. First ed. 2007, Amsterdam, Boston: Butterworth-Heinemann.

21. Springer handbook of materials measurement methods, ed. H. Czichos, T. Saito, and L. Smith. 2006: Springer.

22. IUCr. Crystallographic Information Framework Version 1.1 Working Specification. International Union of Crystallography. 2002 [cited 1 November 2007]; Available from: http://www.iucr.org/iucr-top/cif/spec/version1.1/index.html.

23. Gruber, T.R., Toward Principles for the Design of Ontologies Used for Knowledge Sharing. International Journal Human-Computer Studies, 1995. 43(5-6): p. 907-928.

24. Hunter, J. and K. Cheung. Generating eScience Workflows from Statistical Analysis of Prior Data. in APAC'05. 2005. Royal Pines Resort, Gold Coast.

25. Cheung, K., et al. SCOPE - A Scientific Compound Object Publishing and Editing System. in 3rd International Digital Curation Conference. 2007. Washington DC, USA.

26. Open Archives Initiative Object Reuse and Exchange. [cited 14 Jan 2008]; Available from: http://www.openarchives.org/ore/.

27. Sirin, E., et al., Pellet: A practical OWL-DL reasoner. Web Semantics: Science, Services and Agents on the World Wide Web, 2007. 5(2): p. 51-53.