Origins of The IEEE Standard Upper Ontology

Ian Niles, Adam Pease
Teknowledge
[iniles,apease]@teknowledge.com

Introduction

Research and applications in computer science are creating the need for precise definitions of the concepts that make up our world. Web searching is handicapped by the limitations of specifying search criteria in terms of keywords rather than concepts. Automated natural language understanding, both oral and written, is severely limited by the ambiguity of language. Software engineering is limited by the need for engineers to define concepts to model the world.

Computers exist in a world similar to Europe in the Middle Ages in which tiny principalities each had their own language or dialect. Worse yet, these dialects are impoverished and they enable the computers to say only very specific and limited things. In order to enable continued progress in ecommerce and software integration, we must give computers a common language with a richness that more closely approaches that of human language.

Integrating the meaning (or semantics) of databases and programs is crucial for creating software that is reliable and scalable. The use of ontologies to specify semantics is emerging as a promising technique for software integration. Creators of different components often assume they understand the terms in the same way. The reality is that is rarely the case. Even the best-documented code has implicit assumptions and ambiguity in the definition and usage of terms.

Research in several areas including computer science, artificial intelligence, philosophy, library science and linguistics are helping to meet these needs. All these fields have experience with creating precise and standard descriptions and terminology for the things that make up our world. (Sowa 2000)

Current research is hampered by several issues. Computer scientists and philosophers lack consensus in their communities for creating the very large, wide-coverage ontologies that are needed, although they have the necessary formal languages to do so. Librarians and linguists have the charter to create large ontologies but those ontologies have typically lacked the formal definitions needed for reasoning and decision-making. In fact, it is probably fair to claim that to date no group has taken full advantage of the vast body of historical work in this area One group that has developed a large formal ontology is Cycorp. However, ithas released only a small part to the public, retains proprietary rights to the vast bulk of their ontology, (Lenat 1995) and the contents of the ontology have not been subject to extensive peer review

In this paper we discuss the origins of the Standard Upper ontology effort (SUO, 2001), the strategy used in created the current working draft of the merged ontology, and the current state of the ontology .

The Standard Upper Ontology

Recognizing both the need for large ontologies and the need for an open process leading to a free, public standard, a diverse group of people has come together to make such a standard a reality.

The Standard Upper Ontology (SUO) will be an upper level ontology that provides definitions for general-purpose terms and acts as a foundation for more specific domain ontologies. It is estimated to contain between 1000 and 2500 terms plus roughly ten definitional statements for each term.

Purpose of Project

- The standard will be suitable to support knowledge-based reasoning applications.
- This standard will enable the development of a large (20,000+) general-purpose standard ontology of common concepts, which will provide the basis for middle-level domain ontologies and lower-level application ontologies.
- The ontology will be suitable for "compilation" to more restricted forms such as XML or database schemata. This will enable database developers to define new data elements in terms of a common ontology, and thereby gain some degree of interoperability with other compliant systems.
- Owners of existing systems will be able to map existing data elements just once to a common ontology, and thereby gain a degree of interoperability with other

representations that are compliant with the SUO.

- Domain-specific ontologies that are compliant with the SUO will be able to interoperate (to some degree) by virtue of the shared common terms and definitions.
- Applications of the ontology will include:
  - E-commerce applications from different domains which need to interoperate at both the data and semantic levels.
  - Educational applications in which students learn concepts and relationships directly from, or expressed in terms of, a common ontology. This will also enable a standard record of learning to be kept.
  - Natural language understanding tasks in which a knowledge-based reasoning system uses the ontology to disambiguate natural language terms and structures.

The current working paper for the standard is an ontology in a version of KIF (Genesereth, 1991) called SUO-KIF. It is a somewhat simplified version of KIF that is itself a separate standards effort proposed as a new work item in March 2001 under the ISO.

The SUO email list was created in May of 2000 and quickly had over 150 subscribers. The Project Authorization Request (PAR) which details the scope and purpose for the SUO effort was submitted to the IEEE in October and was approved as working group P1600.1 in December with James Schoening as chair.

Strategy - Reflective Equilibrium

Three common strategies are used in building ontologies, viz. "top-down", "bottom-up", and "middle-out". Each of these strategies has advantages and limitations. A "top-down" approach begins by identifying all of the relevant, philosophico-linguistic concepts, organizing them into a compact, high-level taxonomy and system of axioms, and then proceeding from there to more specific concepts and axioms. The main advantage of the top-down approach is that, if it is executed correctly, it results in a structure which represents a bird's eye view of the world and which should make the task of defining domain-specific content relatively trivial.

The limitation of the "top-down" approach to ontological engineering is a practical one. This approach is embodied in the 2,500-year history of western philosophy, and this discipline has not yielded any conclusions that a majority of philosophers could agree on. It has given us a nice knowledge representation language, the first-order predicate calculus, but it has not resulted in any consensus about the nature of the most general categories of the world and the mind. If some of the deepest thinkers of the western world still haven't come to any solid conclusions after working through these problems for 2,500 years, what hope is there that we mere mortals can do any better? Furthermore, traditional metaphysics and epistemology may not even be relevant to the task of constructing an upper-level ontology. This sounds paradoxical since computer scientists and philosophers use the same word to refer to very similar formal structures of concepts and axioms. However, in this case, appearances are misleading. Ontological philosophers have a radically different aim from ontological engineers. The latter are attempting to construct an engineering artifact that will promote information sharing between heterogeneous automated systems, while the former are trying to get clear about the nature of ultimate reality and what we can know about it. Since the aims in the two cases are so different, much of the theoretical content of philosophy will not be relevant to the task of defining an upper-level engineering ontology. In the context of the SUO project, we're not trying to stake claims on the nature of the world; we're essentially trying to develop the semantic content of a language that can express such claims (and many other claims as well).

Opposed to the so-called "top-down" approach to ontological engineering are the "bottom-up" and "middle-out" approaches. Both of these latter approaches begin with domain-specific concepts and then develop the ontology up or out from there. The advantage of this is that it is, generally speaking, much easier to secure agreement about the sorts of lower level concepts that belong in an ontology. Accordingly, these approaches are very useful for creating domain-specific ontologies from scratch. However, they are not as applicable when the goal is to reuse a stock of variously conceived conceptual content. These approaches do not really lend themselves to the task of merging a collection of very different ontologies. Without some sort of foundation to which the various ontologies can be tied, it is difficult to relate these ontologies to one another.

In constructing the merged ontology, we followed an approach that might, after John Rawls' term, be labeled "reflective equilibrium".

As we see it, this approach has all of the advantages and none of the disadvantages of the "top-down", "middle-out", and "bottom-up" approaches. This novel approach can be described as follows. We begin with a set of high-level categories and then successively incorporate new content under these categories. The inevitable incompatibilities between new content and existing content are resolved on a case-by-case basis. If we find that some high-level categories are not needed to support lower-level nodes, then they are excised from the conceptual structure. If we find that a lower-level structure can be simplified by locating part of its semantic endowment in a new higher-level node (from which the lower-level structure can inherit this content), we can add that node to the ontology. Self-consistent chunks of ontological content sheds some light on higher-level concepts, while explicating the latter concepts partially determines the nature of subsumed concepts. Eventually (if we allow ourselves to be optimistic) we will reach a stage where the demands of both the high-level concepts and the lower-level chunks have been satisfied in the simplest manner possible. At that stage, we will have reached equilibrium. This resulting ontology will not contain all of the original content in the axiomatic chunks or the definitions of high-level concepts, but it should be comprehensive in the sense that it provides a hook on which to hang any domain-specific ontology.

Tactics – Sifting through the Mound

After figuring out how we would approach the task of creating the merged ontology, we proceeded to the next level of detail and outlined a four-step procedure for actually constructing the ontology.

1. Identify relevant, freely available content.
2. Translate content into the same KR language.
3. Merge content into a single, consistent, comprehensive ontology by using "reflective equilibrium".
4. Submit content to SUO mailing list for extension/revision.

The first step was to identify all high-level ontological content that did not have licensing restrictions. This content included the libraries of ontologies available on the Ontolingua server and from ITBM-CNR, John Sowa's upper-level ontology (Sowa, 2000), Russell and Norvig's upper-level ontology (Russell & Norvig, 1995), James Allen's temporal axioms (Allen, 1984), Casati and Varzi's formal theory of holes (Casati & Varzi, 1995), Barry Smith's ontology of boundaries (Smith, 1994, 1996), Nicola Guarino's formal mereotopology (Borgo et al, 1995, 1996), and various formal representations of plans and processes including CPR (Pease, 1997) and PSL (Schlenoff et al, 2000). After all of the relevant content was identified and linked to the SUO web site (http://suo.ieee.org/refs.html), it was translated into a single knowledge representation language. The language used is a simplified variant of KIF as mentioned abouve and a specification of which can be found at http://suo.ieee.org/suo-kif.html.

After the translation of the ontological content into SUO-KIF (known as the "syntactic merge") had been completed, we were faced with the much more difficult task of the "semantic merge", i.e. combining all of the various ontologies into a single, consistent, and comprehensive framework. The ontologies were first divided into two classes, viz. those defining very high-level philosophico-linguistic concepts and those defining lower-level notions. The first class contained John Sowa's upper-level ontology and Russell and Norvig's upper-level ontology, and the second class contained everything else. After this partition was completed, the two upper-level ontologies were melded into a single upper-level ontology. Since both ontologies are very compact and contain a significant amount of overlapping content, this merge did not pose any serious practical problems. This merged upper-level ontology was then used as the foundation for aligning all of the other content that had been converted to the SUO-KIF format.

One of the biggest challenges in merging and aligning ontologies is that many of the chunks that have to be merged and aligned will be, to a lesser or greater extent, incompatible. In some cases, the incompatibilities can be smoothed over by tweaking definitions of concepts or formalizations of axioms; in other cases, wholesale theoretical revision may be required. In hooking the lower-level content into the tip of the ontology, we came across four possible cases.

1. New content that is useful and relatively noncontroversial
2. New content that is either controversial or without an engineering application.
3. Perfect overlap with already existing

content

4. Partial overlap with already existing content

In the first case, nothing in the tip of the ontology corresponds with the concept/axiom to be mapped, and the concept/axiom is deemed to be useful and not to violate any cherished philosophical principles. Once the decision was made to include a concept/axiom in the merged ontology, it was simply a matter of finding a place for it. In some cases, this involved the creation of some intermediate levels of concepts between existing concepts and new content. An example of the first case is the subhierarchy of numbers that was taken from the KIF-Numbers ontology on the Ontolingua server.

```
(subclass-of RealNumber Number)
(subclass-of RationalNumber RealNumber)
(subclass-of PositiveRealNumber
RealNumber)
(subclass-of NegativeRealNumber
RealNumber)
(subclass-of Integer RationalNumber)
(subclass-of EvenInteger Integer)
(subclass-of OddInteger Integer)
(subclass-of NaturalNumber Integer)
(subclass-of NonnegativeInteger Integer)
(subclass-of NegativeInteger Integer)
(subclass-of PositiveInteger Integer)
(subclass-of PositiveNumber Number)
(subclass-of NegativeNumber Number)
(subclass-of ComplexNumber Number)
```

In the second case, the new concept/axiom was judged to be out of place in a schema that we hope will have broad application and acceptance. This sort of judgement is of course somewhat subjective, but this shouldn't diminish the importance of ontology merging/alignment for the same reason that it doesn't dull the significance of legal decisions. An example of the second case is the concept of "Mediating Entity", which appears in John Sowa's upper-level ontology. This concept is derived from the work of the philosopher Charles S. Peirce (Peirce, 1932), and it corresponds to his notion of "Thirdness", i.e. anything that brings two other things into some sort of relationship. Although this notion may be philosophically indispensable, it was difficult to justify its inclusion in an engineering-oriented context, and, for this reason, it was removed from the merged ontology.

In the third case, there is perfect overlap between an element of the merged ontology and the concept/axiom to be mapped - the terms may differ but the new concept has the same semantic content as a concept already in the merged ontology or, with respect to axioms, there is a logical equivalence between the new axiom and an existing axiom. An example of this sort of case occurs with respect to the various mereotopological theories that were referenced in the merged ontology. Some philosophers use 'part-of' as the primitive notion to frame their axioms, some use 'overlaps', and still others use some notion of 'connection'. However, these notions are interdefinable, and thus axioms that are framed in terms of one concept can be easily translated into the terms of the other concepts.

The final case that we encountered in merging the tip of the ontology with the lower-level content is a partial overlap in meaning between the new content and existing concepts or axioms in the merged ontology. An example is the overlap between the concepts of 'Class' and 'Set'. The concept of 'Class' occurs in John Sowa's ontology, where it refers to a collection of items that form something like a natural kind. The concept of 'Set', on the other hand, occurs in the set theory ontology available on the Ontolingua server. To a large extent, these two concepts behave in the same way. Things can be elements of classes and sets, subclasses and subsets are well defined, and classes and sets can be partitioned, disjoint, etc. Despite this strong similarity between the two concepts, it was decided that both of them should be maintained in the merged ontology, because there is a crucial difference in meaning here. While any collection of items qualifies as a set, a class includes only items that share a property or set of properties.

Outstanding Challenges

The major outstanding problem with regard to the merged ontology is inconsistency between engineering-relevant chunks of theoretical content. The clearest example of this sort of clash thus far has been a discussion, on the SUO mailing list, about whether the ontology should have a 3D orientation or a 4D orientation or, as the debate has also been framed, an endurantist perspective or a perdurantist perspective. The 3D orientation assumes that there is a basic, categorial distinction between objects and processes, while the 4D orientation does not. The 3D orientation posits that objects, unlike processes, are completely present at any moment of their existence, while a 4D orientation regards everything as a space-time worm (or a slice of such a worm). On the latter view, paradigmatic processes and objects are merely opposite ends of a continuum of spatio-temporal phenomena. Another example of substantive disagreement concerns the representation of time. As mentioned above, we have aligned James Allen's temporal axioms with the merged ontology. Although these axioms have a clear engineering application, they do assume an interval-based representation of time, which is inconsistent with a point-based representation. Since this latter representation is more applicable in some contexts than the interval-based representation, we apparently need some formal mechanism for accommodating theoretical inconsistency.

The question, then, is how do we handle cases like these, given that our goal is to construct a single, consistent, and comprehensive ontology. It will be unfortunate if we cannot reach this goal, but perhaps it is unattainable. If it is unattainable, then perhaps the best we can do is to make clear the various representational choices and bundle them up in consistent and independent packages and, where possible, state mappings between corresponding packages. One elegant suggestion was (Menzel) to create a lattice from these various modules. The top node of this lattice would be the merged ontology, and each level below the top node would represent inconsistent formal theories that could be used in conjunction with the merged ontology. Thus, each path through the lattice from the top node to a lower-level node would result in a formal theory that is self-consistent, but inconsistent with various other representational choices provided by the ontological lattice.

Another challenge is in making the contents of the ontology accessible and intelligible to the widest audience. Considerable work has been done on tools for browsing ontologies (Swartout et al 1996), (Farquhar et al, 1996) and we expect to make the merged ontology available in a more friendly form that its current incarnation as a very long text file of KIF expressions. We also plan to release versions in different languages which can be generated automatically from the KIF. These additional languages will include DAML (DAML, 2001), XML, and possibly ACE (Fuchs & Schwitter, 1996).

References

Allen, James. F. (1984) "Towards a general theory of action and time," *Artificial Intelligence*, 23, 123-154.

Borgo S., Guarino N., and Masolo C. (1997) "An Ontological Theory of Physical Objects," in L. Ironi (ed.), *Proceedings of Eleventh International Workshop on Qualitative Reasoning (QR'97)*, Cortona (Italia), 3-6 Giugno, 223-231.

Borgo S., Guarino N., and Masolo C. (1996) "A Pointless Theory of Space Based on Strong Connection and Congruence**,"** in L. Carlucci Aiello and J. Doyle (eds.), *Principles of Knowledge Representation and Reasoning (KR96)*, Morgan Kaufmann.

Casati, Roberto and Achille Varzi (1995) *Holes and Other Superficialities*, MIT Press, Cambridge, MA.

Casati, Roberto, Barry Smith and Achille Varzi (1998) "Ontological Tools for Geographic Representation", in Nicola Guarino (ed.), *Formal Ontology in Information Systems*, IOS Press (Frontiers in Artificial Intelligence and Applications), Washington DC, 77–85. http://wings.buffalo.edu/philosophy/faculty/smith/articles/fois(csv).html

DAML (2001). DARPA Agent Markup Language, http://wwww.daml.org/

Farquhar, A., Fikes, R., and Rice, J., (1996). The Ontolingua Server: a Tool for Collaborative Ontology Construction, Proceedings of the 10th Banff Knowledge Acquisition for Knowledge Based System Workshop (KAW95), Banff, Canada, http://ontolingua.stanford.edu/

Fuchs, E. & Schwitter, R. (1996). Attempto Controlled English (ACE). In Proceedings of The First International Workshop On Controlled Language Applications. Katholieke Universiteit Leuven, pages 124-136, Belgium.

Genesereth, M. (1991). Knowledge Interchange Format. In Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR-91), J. Allenet al., (eds), Morgan Kaufman Publishers, 1991, pp 238-249. See also http://logic.stanford.edu/kif/kif.html.

Lenat, D. (1995). "Cyc: A Large-Scale Investment in Knowledge Infrastructure." Communications of the ACM 38, no. 11 (November).

Menzel, C., personnal communication.

Pease & Carrico, (1997), "JTF-ATD Core Plan Representation: A Progress Report". In Proceedings of the AAAI-97 Spring Symposium on Ontological Engineering.

C. S. Peirce. Collected Papers of Charles Sanders Peirce, vol.2, C. Hartshorne, et al: (eds.), Harvard University Press, 1932.

Russell, Stuart and Peter Norvig (1995) *Artificial Intelligence: A Modern Approach*.

Schlenoff, C., Gruninger M., Tissot, F., Valois, J., Lubell, J., Lee, J., (2000). The Process Specification Language (PSL): Overview and Version 1.0 Specification, NISTIR 6459, National Institute of Standards and Technology, Gaithersburg, MD.

Smith, Barry (1994) "Fiat Objects," in N. Guarino, L. Vieu and S. Pribbenow (eds.), *Parts and Wholes: Conceptual Part-Whole Relations and Formal Mereology, 11th European Conference on Artificial Intelligence*, Amsterdam, 8 August 1994, Amsterdam: European Coordinating Committee for Artificial Intelligence, 15–23.

Smith, Barry (1996) "Mereotopology: A Theory of Parts and Boundaries", *Data and Knowledge Engineering*, 20, 287–303.

Sowa, John (2000) *Knowledge Representation*, Brooks/Cole, Pacific Grove, CA.

SUO, (2001), The IEEE Standard Upper Ontology web site, http://suo.ieee.org

Swartout, B., Patil, R., Knight, K. and Russ, T. (1996). Ontosaurus: a tool for browsing and editing ontologies . Gaines, B.R. and Musen, M.A., Ed. Proceedings of Tenth Knowledge Acquisition Workshop. pp.69-1-69-12 (http://ksi.cpsc.ucalgary.ca/KAW/KAW96/swartout/ontosaurus_demo.html)