

Towards a Standard Upper Ontology

Ian Niles, Adam Pease

Teknowledge Corporation

1810 Embarcadero Road, Palo Alto, CA 94303

[iniles,apease]@teknowledge.com

Abstract — The Suggested Upper Merged Ontology (SUMO) is an upper level ontology that has been proposed as a starter document for The Standard Upper Ontology Working Group, an IEEE-sanctioned working group of collaborators from the fields of engineering, philosophy, and information science. The SUMO provides definitions for general-purpose terms and acts as a foundation for more specific domain ontologies. In this paper we outline the strategy used to create the current version of the SUMO, discuss some of the challenges that we faced in constructing the ontology, and describe in detail its most general concepts and the relations between them.

Categories & Descriptors — I.2.4 [Knowledge Representation Formalisms and Methods]: Artificial Intelligence – *representations (procedural and rule-based), semantic networks.*

General Terms — Documentation, Languages, Standard-ization, Theory.

Keywords — Ontologies, Knowledge Interchange Format.

1. Introduction

Cutting-edge software applications are creating the need for a complete set of precise concepts. Web searching is handicapped by the fact that the user must specify his concepts in terms of keywords. Automated natural language understanding, both spoken and written, is limited by the ambiguity of language. Integration of software is hampered by the fact that engineers often define modeling concepts on an ad hoc and informal basis.

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. Unfortunately, there is now, as things stand, a trade-off between precision and expressiveness. On the one hand, computer-readable languages are impoverished - they permit computers to represent only very specific and limited things. On the other hand, human languages can state almost anything anyone would ever want to say. However, so many of the terms and structures of human languages are vague or ambiguous that these languages are not very useful for specifying meanings to a computer.

Research in computer science, artificial intelligence, philosophy, library science, and linguistics are helping to meet the need for a comprehensive, formal ontology. All of these fields have experience with creating standard descriptions and terminology for the entities and

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

FOIS'01, October 17-19, 2001, Ogunquit, Maine, USA.

Copyright 2001 ACM 1-58113-377-4/01/0010...\$5.00.

events that make up our world [18]. However, none of these fields has been able, on its own, to construct a standard, upper-level ontology. 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.

Recognizing both the need for large ontologies and the need for an open process leading to a free, public standard, a diverse group of collaborators from the fields of engineering, philosophy, and information science have come together to make the Standard Upper Ontology (SUO) a reality. 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.

The Standard Upper Ontology (SUO) will provide definitions for general-purpose terms, and it will act as a foundation for more specific domain ontologies. It is estimated that it will eventually contain between 1000 and 2500 terms and roughly ten definitional statements for each term. The SUO will have a variety of purposes, some of which can be glossed as follows.

- Design of new knowledge bases and databases. Developers can craft new knowledge and define new data elements in terms of a common ontology, and thereby gain some degree of interoperability with other compliant systems.
- Reuse/integration of legacy databases. Data elements from existing systems can be mapped just once to a common ontology.
- Integration of domain-specific ontologies. Such ontologies (if they are compliant with the SUO) will be able to interoperate (to some degree) by virtue of shared terms and definitions.

2. SUMO Methodology

The SUMO (Suggested Upper Merged Ontology) is an ontology that was created at Teknowledge Corporation with extensive input from the SUO mailing list, and it has been proposed as a starter document for the SUO Working Group. The SUMO was created by merging publicly available ontological content [12] into a single, comprehensive, and cohesive structure. This content included the ontologies available on the Ontolingua server, John Sowa's upper level ontology, the ontologies developed by ITBM-CNR, and various mereotopological theories, among other sources. The knowledge representation language for the SUMO is a version of KIF (Knowledge Interchange Format) (Genesereth, 1991) called SUO-KIF. This is a somewhat simplified version of KIF, and it is itself a separate proposed standards effort. A specification of the current version of SUO-KIF can be found at <http://suo.ieee.org/suo-kif.html>. The SUMO is a work in progress, and it is growing on a weekly basis. As of July 2001, the ontology contains 654 terms and 2351 assertions. The ontology can be browsed online (<http://ontology.teknowledge.com:8080/rsigma/SKB.jsp>), and source files for all of the versions of the ontology can be downloaded (<http://ontology.teknowledge.com/cgi-bin/cvsweb.cgi/SUO/>).

The procedure that we followed in creating the SUMO can be glossed as follows. 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 [19], Russell and Norvig's upper-level ontology [15], James Allen's temporal axioms [1], Casati and Varzi's formal theory of holes

[4], Barry Smith's ontology of boundaries [17, 18], Nicola Guarino's formal mereotopology [2,3], and various formal representations of plans and processes including CPR [13] and PSL [16]. 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 SUO-KIF.

After the translation of the ontological content (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 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 conceptual structure. Since both of the source ontologies are very compact and contain a significant amount of overlapping content, this merge did not pose any serious practical or theoretical problems. This merged ontology was then used as the foundation for aligning all of the other content that had been converted into SUO-KIF.

We encountered four sorts of cases in aligning lower-level content with the foundational ontology. In the first case, nothing in the tip of the ontology corresponded with the concept/axiom to be mapped, and the concept/axiom was 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.

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 judgment 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 [14], 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.

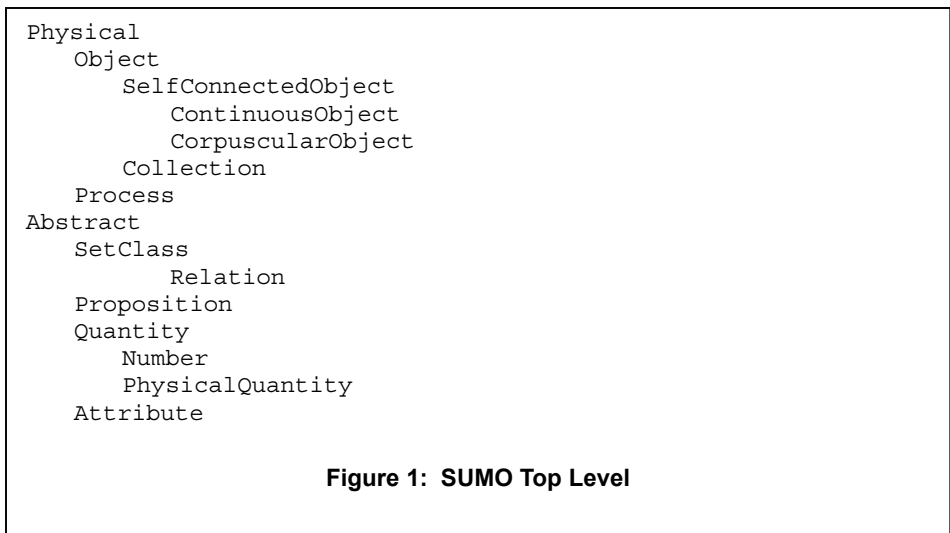
In the third case, there is perfect overlap between an element of the merged ontology and the concept/axiom to be mapped - the terms may differ but the new concept has the same semantic content as a concept already in the merged ontology or, with respect to axioms, there is a logical equivalence between the new axiom and an existing axiom. An example of this sort of case occurs with respect to the various mereotopological theories. 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 with one concept can be easily translated into the other concepts.

The final case that we encountered in merging the foundational 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. This case represents one of the biggest challenges in ontological merging/alignment. Many of the chunks that have to be incorporated 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. Consider, for example, the overlap between the concepts of 'Class' and 'Set'. The concept of 'Class' occurs in John Sowa's ontology,

where it refers to a set 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 SUMO. We will discuss in the next section the precise definitions of these concepts and the relation between them.

3. SUMO Top Level

Although the procedure and examples outlined in the previous section should be helpful in understanding the ontology, the best way of explaining the structure and content of the SUMO is to systematically present the highest level concepts and the relations between them. The diagram in Figure 1 is a snapshot of these concepts, and the indentations indicate subsumption relations. The root node of the SUMO is, as in many ontologies, ‘Entity’, and this concept subsumes ‘Physical’ and ‘Abstract’. The former category includes everything that has a position in space/time, and the latter category includes everything else.



Under the concept of ‘Physical’, we have the disjoint concepts of ‘Object’ and ‘Process’. The existence and nature of the distinction between these two notions was the subject of much heated debate on the SUO mailing list. According to those who adopt a 3D orientation (or “endurantists”, as they are sometimes called), there is a basic, categorical distinction between objects and processes. According to those who adopt a 4D orientation (the “perdurantists”), on the other hand, there is no such distinction. 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. The current version of the SUMO embodies a 3D orientation by making ‘Object’ and ‘Process’ disjoint siblings of the parent node ‘Physical’. The reason we adopted this orientation is that we wanted to incorporate content from process-related ontologies like PSL and from formal mereotologies (which assume a substantive notion of object).

Immediately under the concept of ‘Object’, there are two disjoint concepts, viz. ‘SelfConnectedObject’ and ‘Collection’. A ‘SelfConnectedObject’ is any ‘Object’ whose parts are all mediately or immediately connected with one another. This definition has been formalized in various ways, and we include the following axiom in the SUMO.

```
(<=>
  (instance-of ?OBJ SelfConnectedObject)
  (forall (?PART1 ?PART2)
    (=>
      (and
        (part-of ?PART1 ?OBJ)
        (part-of ?PART2 ?OBJ))
      (connected ?PART1 ?PART2))))
```

Where `connected` is a generalized notion of connection accounting for mediate and immediate unifying relations over particular divisions of a whole, as in Simons [17], e.g.:

```
(=> (and
  (instance-of ?P1 ?DIV)
  (instance-of ?P2 ?DIV)
  (or
    (immediately-connected ?P1 ?P2)
    (exists (?P3)
      (and
        (instance-of ?P3 ?DIV)
        (connected ?P1 ?P3)
        (immediately-connected ?P2 ?P3))))))
  (connected ?P1 ?P2))
```

The concept of ‘SelfConnectedObject’ is partitioned into two concepts: ‘ContinuousObject’ and ‘CorpuscularObject’. Both of these concepts are borrowed from John Sowa’s ontology. A ‘ContinuousObject’ is an ‘Object’ in which every part is similar to every other in every relevant respect. More precisely, something is a ‘ContinuousObject’ when all of its parts (down to an unspecified level of granularity) have the properties of the whole. Thus, substances like water and clay would be subclasses of ‘ContinuousObject’, as would topographic locations like surfaces and geographic areas. Formally, the distinction between ‘ContinuousObject’ and ‘CorpuscularObject’ can be described with the following axioms in the SUMO.

```
(=>
  (and
    (subclass-of ?OBJECTTYPE ContinuousObject)
    (instance-of ?OBJECT ?OBJECTTYPE)
    (part-of ?PART ?OBJECT))
  (instance-of ?PART ?OBJECTTYPE))

(equal CorpuscularObject (ComplementFn ContinuousObject))
```

Disjoint from ‘SelfConnectedObjects’ are ‘Collections’. ‘Collections’ consist of disconnected parts, and the relation between these parts and their corresponding ‘Collection’ is known as ‘member’ in the SUMO. The following axiom in the SUMO requires that there be no empty collections.

```
(=>
  (instance-of ?COLL Collection)
  (exists (?OBJ)
    (member ?OBJ ?COLL)))
```

Note that this ‘member’ predicate is different from the ‘instance’ and ‘element’ predicates, which relate things to the ‘Classes’ or ‘Sets’ to which they belong. Unlike ‘Classes’ and ‘Sets’, ‘Collections’ have a position in space-time, and ‘members’ can be added and subtracted without thereby changing the identity of the ‘Collection’. Some examples of ‘Collections’ are toolkits, football teams, and flocks of sheep.

Going back up the hierarchy of the ontology, we return to the concept of ‘Physical’ and its other child ‘Process’. In this paper, we will not discuss any of the concepts subsumed by ‘Process’. We will just point out that this was and is continuing to be one of the most challenging parts of the SUMO. We wanted a complete typology of the various sorts of processes, but we found that there were very few resources to guide us here. Although we have incorporated the core of PSL (Process Specification Language) into the SUMO, this formal theory does not attempt to break down the impressive variety of processes. We finally turned to a well-received work that, among other things, attempts to classify over 3,000 English verbs into “semantically coherent verb classes” [11]. There are 48 main verb classes outlined in this work. However, not all of these are relevant to the task of developing a typology of processes. Some of the verb classes relate to static predicates in the ontology rather than to processes, and some classes seem to be syntactically motivated, e.g. the class of verbs that take predicative complements. We have eliminated the verb classes that do not seem to refer to genuine processes, adjusted the hierarchy somewhat, and developed formal axioms for the remaining concepts. We will present this subontology of processes in a future paper.

Returning to the highest level distinction in the SUMO, that between ‘Physical’ and ‘Abstract’, let us now consider the ‘Abstract’ branch. The class ‘Abstract’ subsumes four disjoint concepts: ‘Set’, ‘Proposition’, ‘Quantity’, and ‘Attribute’. ‘Set’ is the ordinary set-theoretic notion, and it subsumes ‘Class’, which, in turn, subsumes ‘Relation’. A ‘Class’ is understood as a ‘Set’ with a property or conjunction of properties that constitute the conditions for membership in the ‘Class’, and a ‘Relation’ is a ‘Class’ of ordered tuples. Note that ‘Relation’ is immediately subsumed by ‘Class’, rather than ‘Set’, because we restrict ‘Relations’ to those ordered tuples that express intensional content. For example, the predicate ‘mother’ in the SUMO does not just denote the set of all ordered pairs whose first element is a mother and whose second element is an offspring of the first element. This predicate also expresses the property or conjunction of properties that comprise the notion of biological motherhood.

The concept of ‘Proposition’ corresponds to the notion of semantic or informational content. However, the SUMO places no size restrictions on this content. Although some ‘Propositions’ are expressed by single sentences, other ‘Propositions’ are expressed by entire books or even libraries of books. This is a broader notion than is used in many ontologies, but it does not seem to be possible to make a principled distinction between the abstract content expressed by one sentence and the abstract content expressed by larger units of discourse. Some examples of ‘Propositions’ would be the story line conveyed by a novel and the musical content denoted by a printed score.

The class of ‘Attributes’ includes all qualities, properties, etc. that are not reified as ‘Objects’. For example, rather than dividing the class of ‘Animals’ under ‘Objects’ into “FemaleAnimals” and “MaleAnimals”, we make ‘Female’ and ‘Male’ instances of ‘BiologicalAttribute’, which is a subclass of ‘Attribute’.

Finally, ‘Quantity’ under ‘Abstract’ is divided into ‘Number’ and ‘PhysicalQuantity’. The former is understood as a count independent of an implied or explicit measurement system, and the latter is taken to be a complex consisting of a ‘Number’ and a particular unit of measure. Thus, 1 meter and 39.37 inches would be two distinct instances of ‘PhysicalQuantity’ in the SUMO, though they can be shown to be equivalent via two conversion axioms in the ontology.

5. Challenges to SUMO

The major problem we encountered in constructing the SUMO was inconsistency between engineering-relevant chunks of theoretical content. The clearest example of this sort of clash has been the debate about whether the SUO should have a 3D orientation or a 4D orientation, as discussed earlier. 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. Chris Menzel made the elegant suggestion that a lattice be created from these various modules. The top node of this lattice would be the SUMO, and each level below the top node would represent inconsistent formal theories that could be used in conjunction with the SUMO. 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.

Inevitably people ask why we would engage in constructing an upper-level ontology, when such an ontology already exists. The Austin-based company Cycorp has devoted fifteen years to creating an ontology that has been used in a wide range of applications. The Cyc ontology is extremely impressive, but there are problems with its use as a standard. Cycorp has released only a small part of its ontology to the public, the company retains proprietary rights to the vast bulk of its ontology [10], and the contents of the ontology have not been subject to extensive peer review. Aside from these limitations of Cyc, there are some distinct advantages of the SUMO. First, the SUMO is the working paper of an IEEE-sponsored open-source standards effort. This means that users of the ontology can be more confident that the ontology will eventually be embraced by a large class of users. Second, the SUMO was constructed with reference to very pragmatic principles. Any distinctions of strictly philosophical interest have been removed from the ontology. This has resulted in a KB that should be simpler to use than Cyc.

References

- [1] Allen, J. (1984) "Towards a general theory of action and time," *Artificial Intelligence*, 23, 123-154.
- [2] 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.
- [3] Borgo, S., Guarino, N., and Masolo, C. (1996) "A Pointless Theory of Space Based on Strong Connection and Congruence," in L. C. Aiello and J. Doyle (eds.), *Principles of Knowledge Representation and Reasoning (KR96)*, Morgan Kaufmann.
- [4] Casati, R. and Varzi, A. (1995) *Holes and Other Superficialities*, MIT Press, Cambridge, MA.
- [5] Casati, R., Smith, B. and Varzi, A. (1998) "Ontological Tools for Geographic Representation", in N. 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](http://wings.buffalo.edu/philosophy/faculty/smith/articles/fois(csv).html)
- [6] DAML (2001). DARPA Agent Markup Language, <http://www.daml.org/>

- [7] 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/>
- [8] Fuchs, E. and Schwitler, R. (1996) "Attempto Controlled English (ACE)," *Proceedings of The First International Workshop On Controlled Language Applications*, Katholieke Universiteit Leuven, Belgium, 124-136.
- [9] Genesereth, M. (1991). "Knowledge Interchange Format," *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR-91)*, J. Allen et al., (eds), Morgan Kaufman Publishers, 1991, 238-249. See also <http://logic.stanford.edu/kif/kif.html>.
- [10] Lenat, D. (1995) "Cyc: A Large-Scale Investment in Knowledge Infrastructure," *Communications of the ACM*, 38, no. 11 (November).
- [11] Levin, B. (1993) *English Verb Classes and Alternations: A Preliminary Investigation*, University of Chicago Press.
- [12] Niles, I and Pease, A. (2001) "Origins of the IEEE Standard Upper Ontology," *Working Notes of the IJCAI-2001 Workshop on the IEEE Standard Upper Ontology*, Seattle, 37-42.
- [13] Pease, A & Carrico, (1997), "JTF-ATD Core Plan Representation: A Progress Report," *Proceedings of the AAAI-97 Spring Symposium on Ontological Engineering*.
- [14] Peirce, C. S. (1932) *Collected Papers of Charles Sanders Peirce*, vol.2, C. Hartshorne, et al (eds.), Harvard University Press.
- [15] Russell, S. and Norvig, P. (1995) *Artificial Intelligence: A Modern Approach*.
- [16] Schlenoff, C., Gruninger M., Tissot, F., Valois, J., Lubell, J., and Lee, J. (2000) *The Process Specification Language (PSL): Overview and Version 1.0 Specification*, NISTIR 6459, National Institute of Standards and Technology, Gaithersburg, MD.
- [17] Simons, P. (1987) *Parts: a Study in Ontology*. Clarendon Press, Oxford.
- [18] Smith, B. (1994) "Fiat Objects," 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, 15-23.
- [19] Smith, B. (1996) "Mereotopology: A Theory of Parts and Boundaries", *Data and Knowledge Engineering*, 20, 287-303.
- [20] Sowa, J. (2000) *Knowledge Representation*, Brooks/Cole, Pacific Grove, CA.
- [21] SUO, (2001), The IEEE Standard Upper Ontology web site, <http://suo.ieee.org>
- [22] 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. (eds.), *Proceedings of Tenth Knowledge Acquisition Workshop*. (http://ksi.cpsc.ucalgary.ca/KAW/KAW96/swartout/ontosaurus_demo.html)