Examining SNOMED from the Perspective of Formal Ontological Principles: Some Preliminary Analysis and Observations

Kent A. Spackman, MD PhD[a], Guillermo Reynoso, MD MBA[b]

[a]*Department of Medical Informatics and Clinical Epidemiology, Oregon Health & Science University, Portland, Oregon, USA.*

[b]*InTerm Medical Terminology Research Center, Foundation for Healthcare Informatics Research,Buenos Aires, Argentina.*

Abstract

The Systematized Nomenclature of Medicine – Clinical Terms (SNOMED CT) is a terminological resource designed to support electronic applications in health and medicine. Its design has evolved over a period of more than thirty years, and continues to evolve. Recently several authors working on formal ontological theory have observed that applying certain principles and constraints to terminology construction may result in a more consistent and useful terminology. In this paper we report on a preliminary analysis of SNOMED CT by two of its developers, from the perspective of a few such formal ontological principles, giving examples of prior design decisions that appear to be supported by these principles as well as examples of prior design decisions that may be at variance with them. We believe that design changes suggested by formal ontological principles have great potential for improving consistency. Empirical evidence of usefulness should accompany theoretically-inspired moves towards more fine-tuned representations of reality.

Introduction

The Systematized Nomenclature of Medicine – Clinical Terms (SNOMED CT) continues to evolve with a goal of being both theoretically well-founded and clinically useful. Recently, several authors have observed that applying formal ontological principles and constraints to terminology construction may result in a more consistent and useful terminology. In this paper we report on a preliminary analysis of SNOMED CT from the perspective of a few such formal ontological principles, giving examples of prior design decisions that appear to be supported by these principles as well as examples of prior design decisions that may be at variance with them. We believe that design changes suggested by formal ontological principles have great potential for improving consistency. We also agree with Welty and Guarino[1] that these changes can result in creation of additional concepts with the same or nearly the same *term*, and some users may tend to view this as duplicative and redundant. It is not possible to be sure *a priori* that all such design changes will improve the terminology's value to its users, and we believe, at least for SNOMED, that empirical evidence of usefulness should be sought to accompany theoretically-inspired moves towards more fine-tuned representations of reality.

The ability of SNOMED CT to scale as a global terminology to be used in heterogeneous scenarios depends on several key factors. A significant one is a documented concept model that enables users to use formalized methods for the development of local extensions or for effectively contributing feedback for collaborative refinement of the terminology. While the concept model underlying the development of SNOMED CT by merging SNOMED RT and CTV3 followed a set of design and modeling principles described elsewhere, those principles may need to be reconsidered in terms of recent advances and experience in the application of formal ontological analysis methodologies that facilitate the explication of the modeler assumptions and ontological decisions.

SNOMED Background

SNOMED Clinical Terms is the latest in a long series of works of terminology developed and distributed by the College of American Pathologists (CAP) for the purpose of encoding, storing, and retrieving information on disease and health. Beginning with the Systematized Nomenclature of Pathology (SNOP) in 1965, and continuing through expansion to the Systematized Nomenclature of Medicine (SNOMED) in 1976 and subsequent major editions in 1979 and 1993, the CAP focused on making a practical and comprehensive terminology that could be used by manual coders as well as by computerized information systems. Beginning in the mid-1990's, the CAP embarked on a radical re-engineering of SNOMED with the understanding that manual coding would become a thing of the past, and that substantial changes were required to support increasingly sophisticated electronic systems in healthcare and public health. As a consequence of this re-engineering and substantial re-work, in cooperation with the Kaiser Permanente "Convergent Medical Terminology" (CMT) project, CAP published the SNOMED Reference Terminology (RT) in 2000[2]. An even larger transformation (more than doubling in size, expansion of the concept model and other features) occurred as a result of merging SNOMED RT with the UK National Health Service's (NHS) Clinical Terms version 3 (CTV3), resulting in the first release of SNOMED Clinical Terms (CT) in January of 2002. Since that time there have been an additional four releases, one every six months. In 2003, the US Government licensed SNOMED CT and the National Committee on Vital and Health Statistics (NCVHS) recommended it as the general terminology for patient medical record information in the US. In the UK, SNOMED CT is a draft national standard and a key element of the NHS National Program for IT. Thus SNOMED is not a theoretical academic exercise, but is being developed with serious expectations and demands for practical usability.

Purpose of SNOMED Clinical Terms

SNOMED Clinical Terms is a terminological resource designed to be implemented in software applications to represent clinically relevant information reliably and reproducibly. Through the use of this information, SNOMED CT enabled applications can support effective delivery of high quality healthcare to individual people and populations. SNOMED CT is an international, multilingual terminological resource that can also represent concepts and terms unique to particular organizations or localities.[3]

Guiding Principles of SNOMED Development and Maintenance

Ever since the 1960's, the College of American Pathologists has regarded coding and classification systems as a vital interest. In 2003, working together with their colleagues from the UK NHS, they reiterated their commitment and outlined several basic principles upon which ongoing work by CAP is premised.[4] Prominent among these principles are commitments to 1) clinical integrity and quality, 2) usefulness for support of patient care, patient safety, audit, research, analysis, and planning, 3) scientific validation, 4) sustainability, with direct input from stakeholders, 5) widespread adoption, 6) protection of legacy data, and 7) accommodation of local needs. These are all laudable and necessary commitments, but in reality there are many constraints on any organization's ability to approach perfection in all of these areas, and there are natural tensions between these principles that require pragmatic and ongoing tradeoffs and judgments. This balancing process is analogous to attempting to find a suitable path towards the optimum in a large but constrained search space. There are natural tensions between sustainability, requiring a significant ongoing commitment of resources, versus widespread adoption, requiring minimal barriers and therefore free access. Government support is the preferred means of resolving this tension. There are also natural tensions between clinical integrity/quality/validity, requiring a significant degree of complexity with ongoing changes (enhancements, it is hoped), versus widespread use with protection of legacy data, requiring simplicity, face validity, and careful attention to backwards compatibility. It is in the context of this tension that analyses based on formal ontological principles must be placed, since one cycle's new formalisms, full of promise to "clean up" our problems, may become the next cycle's follies. SNOMED is demonstrably in this for the long haul, so changes will require due deliberation.

Evolutionary Design

Clinical terminology is difficult, and it is unreasonable to expect it ever to be perfect.[5] Rather than an excuse for ignoring problems in the terminology, this is a recognition that the design must adapt and change in order to continue to serve the needs for which it is intended. Campbell's influential work provided the basic evolutionary design principles upon which SNOMED development is still based.[6] There are six main points:

1. Evolution without pre-ordained design
2. Accumulation of design
3. Heterogeneity
4. Participatory consensus-based approach
5. Semantics-based concurrency control
6. Configuration management

As SNOMED development has continued, these broad principles have been operationalized using three fundamental criteria, abbreviated as "URU". The initials stand for understandability, reproducibility, and usefulness. The first criterion, understandability, makes reference to whether a concept (or other design feature of the terminology) can be fully and unambiguously comprehended by users of the terminology. Understandability is tested by checking to see whether users believe they can tell whether the concept is relevant or not relevant to a given patient or situation. It is clear from this subjective test of understandability that two individuals may believe they understand what is meant, but their understanding may differ significantly. This leads to the need for the second criterion: Reproducibility indicates whether multiple users apply the concept to the same situations. Tests of reproducibility generally depend on independent modeling or coding followed by comparison. Finally, usefulness refers to the level of helpfulness and appropriateness conveyed in a concept or feature. A challenge for clinical terminologies is the need to provide explanation to naive users in order to make a sophisticated and complicated terminology accessible and useful.

The description logic definitions used to classify SNOMED CT support conjunction, existential restrictions, role hierarchies and the SNOMED CT notion of role groups, which can be represented using existential role restrictions in any description logic (DL) language.[7] This set of concept constructors is a small subset of DL features compared with the expressivity of $\mathcal{ALC}$, $\mathcal{SHIQ}$, and others. Future significant changes in the concept model might depend on the support of concept constructors like disjunction, negation and transitive properties. Classification tests have shown that the supertype and subtype relationships inferred by any correct and complete classifier will match those obtained and distributed in the SNOMED tables.

Basic Definitions and Ontological Principles

Here we briefly review some of the definitions and principles that have been proposed for subjecting terminologies to formal analysis. Guarino and Welty have proposed a set of principles collectively known as the OntoClean methodology.[8,9,10] This methodology appears to be gaining acceptance as a guidance and evaluation framework. Fundamental to the method is the idea of a *property*, a term roughly corresponding to what is ordinarily called a *concept* in taxonomies and description-logic based terminologies like SNOMED. From an ontological perspective, SNOMED's concepts such as *disorder*, *substance* and *organism* might be called *properties*. To quote Guarino and Welty, "In this paper we show how a formal ontology of unary properties (corresponding to concepts in taxonomies) can help using the subsumption relation in a disciplined way."[8] In this view, we distinguish properties like *organism* from the real-world bearers of those properties (actual organisms).

OntoClean provides formal definitions of meta-properties, which are a group of special properties characterizing other properties. These meta-properties (see examples below) help in the explication of the intended meaning of concepts from a formal ontological point of view. The ability to derive constraints on subsumption from the value assigned to these meta-properties provides assistance in the evaluation of modeling decisions.

Meta-properties

Meta-properties define characteristics of properties by saying what is or is not necessarily true of the instances of those properties. Here we restate the definitions of four of the meta-properties and provide examples attempting to convey an intuitive understanding of what is meant; readers should refer to the primary sources in the references for formal definitions.[8,9,10]

Rigidity: Guarino and colleagues define rigidity as a property that necessarily holds for all its instances in any instant of time and in every possible world. For example, *dog* is a rigid property because all instances of dog must always be dogs; they cannot be a dog at one time and not a dog at another. On the other hand, *pet* is called anti-rigid, meaning that instances of pet are not necessarily pets, since they could cease being a pet when, for example, they no longer have an owner. This is assuming that what we mean by *pet* makes it dependent on being owned, so a pet dog that no longer has an owner is a stray, not a pet, but it must remain a dog.

Identity: This meta-property aims to characterize what is unique for an entity that allows it to be identified, or re-identified, in different times and places. A property is said to carry an identity criterion if all its instances can be re-identified by a criterion that judges sameness. For example, the property *organism* is said to carry an identity criterion, since any instance of organism can be identified as being the same across time, based on biological criteria. On the other hand, the property *asymmetric* would be said not to have an identity criterion, since it is not possible to define criteria to determine whether two instances of asymmetry are necessarily the same.

Unity: To hold the *unity* meta-property, every instance of a property must be an intrinsic whole. The determination of wholeness can depend on topological wholeness, or, alternatively, on a morphological, functional, or other relation. The relation that determines that a property carries a unity condition is called its equivalence relation. A property is said to have *anti-unity* if all its instances are not intrinsic wholes. For example, the property *water* has anti-unity, because there is no sense in which one can specify a relation that determines that instances of water are whole. In contrast, the property *lake* (which consists of water but is not an instance or subtype of water) can have a topological relation that defines its whole (based on the boundaries of the lake bed and the surface of the lake), and can therefore carry the unity meta-property.

Dependence: The meta-property *dependence* implies that all the instances of a given property require the existence of some instance of another entity that is not part of the former. For example, the property *mother* requires the existence of a child (at some point in time), and therefore is dependent. In contrast, the property *female* is independent.

Continuants and Occurrents

In addition to the meta-properties defined in the OntoClean methodology, we believe the distinction between continuants and occurrents, as defined, for example, by Smith and colleagues, provides potentially valuable insights for structuring clinical terminology resources like SNOMED CT.[11,12,13,14] A continuant is an entity that has no temporal parts, and therefore can be understood to exist in a slice of time. Objects, persons, substances, and locations are all in this category. On the other hand, occurrents have temporal duration. Procedures, processes and movements fit into this category.

Taxonomic constraints

The value of distinguishing OntoClean's meta-properties, and the fundamental properties of continuant and occurrent, is that these provide perspectives that enable us to eliminate inconsistencies in terminology hierarchies based on subsumption constraints. In other words, the is-a relation should behave in a consistent manner, and these constraints help us to identify possible inconsistencies and eliminate them.

The following seven constraints represent merely a subset of all possible constraints that might be generated by formal ontological analysis. However, these appear to be important and potentially very useful. If we let "x $\not\subset$ y" mean that properties (concepts) having meta-property x should not be subsumed by (should not have an "is-a" relationship to) any property having meta-property y, then the first five in this list are restatements of Guarino's constraints,[8] and the last two are restatements of, for example, constraints expressed by Fielding[11] and Smith[13].

1. Unity $\not\subset$ Anti-unity
2. Non-unity $\not\subset$ Unity
3. Rigidity $\not\subset$ Anti-rigidity
4. Non-identity $\not\subset$ Identity
5. Independent $\not\subset$ Dependent
6. Continuant $\not\subset$ Occurrent
7. Occurrent $\not\subset$ Continuant

Examples of Taxonomic Constraints

<u>Example concepts</u>: aspirin (product), aspirin (substance)
<u>Constraint</u>: Unity $\not\subset$ Anti-unity

Aspirin (acetylsalicylic acid, ASA) is used to name an ingredient and also to name a class of prepared product that contains ASA. It is an exemplar of a systematic decision in SNOMED to separate ingredient substances from the products of which they are made, even though they have the same name. Ingredient substances would be properties with anti-unity, but the products of which they are made would be properties with unity. Thus formal ontological principles confirm our decision to separate them, but there is evidence that not everyone agrees with the decision. In particular, the editors of the US National Library of Medicine's Unified Medical Language System (UMLS) Metathesarus, which incorporates SNOMED CT into its structure, have decided to leave these two SNOMED codes (aspirin product, aspirin substance) linked to the same concept unique identifier (CUI), and likewise for all other product/substance pairs in SNOMED. (B. Humphreys, personal communication). In other words, they are representing a concept "aspirin" that does not differentiate between the drug product itself and the stuff of which it is made. We also considered this approach because it initially appeared simpler (one concept instead of two), but eventually rejected it because of the difficulties it creates in correctly representing subsumption hierarchies of drug ingredients and drug products. We recognized that we have to accept, and explain to our users, that there will be two concepts carrying the simple name *aspirin* and they will have to choose between them because they are truly different. We provide the (substance) and (product) phrases in the fully specified name to help users to see the difference.

This example nicely illustrates the kinds of tensions that sometimes arise between the requirements of formal rigor and the (apparent) requirements of commonsense thinking and simplicity.

Parenthetically, this mismatch between SNOMED and UMLS once again confirms the interpretation of the UMLS CUI, proposed by Campbell et al,[15] that claims that it must be viewed as representing extensional meaning since it will not always match the intensional meaning of its source vocabularies.

<u>Example concepts</u>: infectious agent, bacterium
<u>Constraints</u>: Rigid $\not\subset$ Anti-rigid, Independent $\not\subset$ Dependent

The property *infectious agent* would be called a *role* in OntoClean. It is an anti-rigid property (in some possible world, all instances can possibly be non-infectious) and dependent on an infectious relationship between the agent and an infected (or perhaps infectable) organism. The property *bacterium*, on the other hand, would be called a *type*. It is a rigid property, carrying identity, and independent. *Infectious agent* currently subsumes *bacterium* in SNOMED. The taxonomic constraints suggest this is inconsistent and should be changed. Once again, there is tension between simplicity and usability on the one hand and formal rigor on the other, since practical use calls for a simple categorization of infectious agents, and the simplest solution appears to be an is-a relationship from bacteria, fungi, parasites, viruses and prions to infectious agent. However, we agree that this role vs. type distinction provides a useful criterion to untangle the taxonomic primitive backbone which, as noted by Welty and Guarino[1] should consist only of rigid properties, although strict adherence to this idealized structure may not always be possible.

This distinction also helps to prioritize the incorporation of new attributes into the SNOMED CT concept model. Adding new attributes results in a more faithful representation of meaning and avoids inconsistencies. As another example of the use of attributes to eliminate incorrect is-a's, consider the relationship of *insulin*, *hormone*, and *antidiabetic agent*, as in Alan Rector's tutorial:[16]

$$\text{Insulin} \subseteq (\text{Antidiabetic agent})$$
$$\text{Antidiabetic agent} \subseteq (\text{protein} \cap \exists\text{hasFunction.HormonalAction})$$

Could instead be modeled as:

$$\text{Insulin} \subseteq (\text{protein} \cap \exists\text{hasFunction.AntidiabeticAction})$$
$$\text{Hormone} = (\text{substance} \cap \exists\text{hasFunction.HormonalAction})$$

Certain semantic categories like *physical object*, *social context*, *substance* and *organism* may benefit from this kind of analysis focused on the type/role distinction in the upper level hierarchy.

Other semantic categories like *finding*, *disorder* and *procedure* may realize less benefit from this methodology, and therefore the case for applying it should be based on future research.

<u>Example concepts</u>: morphologic abnormality, pathological process.
<u>Constraint</u>: Continuant $\not\subset$ Occurrent

Early work on SNOMED RT involved significant discussion and consensus-building resulting from dual independent modeling (dissection, definition) of concepts, followed by examination of differences, as a means of seeking reproducibility. An early disagreement arose about *acute inflammation*. From a clinical examination perspective this term described the combination of redness, pain, swelling and heat of an inflammatory process. From a histological perspective this term described the existence of an infiltrate of acute inflammatory cells. Although described using the same words, the two concepts are very different. If we assume one meaning refers to an acute inflammatory process, an occurrent, and the other meaning refers to an acute inflammatory infiltrate, a continuant, then it is clear we need two different codes and that neither can subsume the other. In fact, they should be in totally different hierarchies: the structure in the morphologic abnormality hierarchy, and the process in the pathological process hierarchy. It is instructive to realize (and useful to apply as a general rule) that the process – structure distinction provided by the words *infiltration* and *infiltrate* may not be provided by the words used commonly for other situations. In this example, we speak of *inflammation* but not "inflammate"; instead we use the same word for both meanings. This re-emphasized for us a universally known but frequently forgotten lesson that simply using the same words, even technically detailed words, is no guarantee of meaning the same thing.

<u>Example concepts</u>: morphologic abnormality, disorder
<u>Constraint</u>: Continuant $\not\subset$ Occurrent

Early in the Kaiser CMT project, there was a discussion about whether SNOMED III morphology (M) codes and SNOMED III disease (D) codes meant the same thing. There was consensus that there is a clear difference between the structural "snapshot" of a disease that is observed by a pathologist in a tissue section, and the temporally extended disease that is experienced by the patient. This decision accounts for the apparent duplication of codes that have very similar names, one for morphologic abnormality and the other for disorder. It seems fairly easy to explain that we have both "neoplasm, benign (morphologic abnormality)" and "benign neoplastic disease (disorder)". But it is more difficult, to some, to follow the same logic to "Burkitt lymphoma (morphologic abnormality)" and "Burkitt's lymphoma (disorder)". Commonsense thinking seems to lead people to believe that we should have only one code for Burkitt's lymphoma.

Singular versus Plural Naming

Formal ontologies carefully distinguish between single whole entities and groups of wholes. Typical thesaurus construction tends to use plurals to describe more general classes, in order to signal to users that the code represents a category. SNOMED III (1993) used plurals in this way for "headers", which were published in uppercase and carried a data field that set them apart (Eclass=00). CTV3 routinely used plurals for higher-level categories and singular tense at lower levels. The transition to SNOMED CT was accompanied by a systematic effort, unfortunately still incomplete, to convert these names to singular tense unless the intended meaning actually implied multiples. For example, "infiltrations (procedure)" is the general procedure subsuming procedures such as "intradermal infiltration of steroid (procedure)". "Infiltration (procedure)" should be its name. Although this decision helps SNOMED align better with an ontological rigor, some users have told us they would like us to present the hierarchies using plurals for the upper level categories, because they feel it would look better. This may result from a mental habit of using hierarchies to name a set of things, and then name the things in the set.

Beyond the necessary attention to singular-plural naming, there is additional attention required to differentiate wholes from collections.

Anatomy and Mereology

Although beyond the scope of this paper, we also want to mention the anatomy model that SNOMED has used to support classification of findings, disorders and procedures. The use of structure-entire-part (SEP) triplets to represent anatomy was inspired by the challenges of combining SNOMED RT and CTV3. More theoretically-oriented justification of the SEP model were independently developed by Schulz et al.[17] This pragmatically-oriented model of body structures has been sufficient for the great majority of concept definitions. While the SEP triplet implementation present in SNOMED CT has been a significant improvement to the SNOMED RT model, certain aspects like the relationship between microscopic - macroscopic structures and part – region modeling would require further analysis. The mereology foundations have a significant impact in how semantic categories are structured by the classifier, and therefore merit further research.

Some Observations and Discussion

In a large terminology such as SNOMED (over 350,000 concepts), there are bound to be errors and inconsistencies from the perspective of formal ontology. In order to properly understand their source and determine what to do about them, it is necessary to know whether they are attributable to design decisions and therefore intentional, or are unintentional errors that, with ongoing maintenance, are being eliminated without the need for further design work. Attempting to determine these questions by examining the terminology alone without asking the developers about the current state of the design, is like reading tea leaves: it is highly subjective, unreliable, and prone to misinterpretation.

OntoClean appears to bring good organizing principles to the modelling of several of SNOMED's semantic categories such as physical object, social context, organism, and substance. In general, untangling these hierarchies is very desirable. The applicability of some of the formal ontology principles in providing consistent guidance in the very large areas like clinical finding and procedure is not as clear, and appears to require further elaboration of ontological foundations as well as further study of the impact of newly proposed distinctions on the structure of the terminology. Yet these represent the most numerous and important of SNOMED's content. Since ontological analysis is labor-intensive, implementing its constraints in certain areas of SNOMED or adding modelling criteria should be based on proven benefits and/or pilot experiences. Meeting the needs of users need not be in conflict with detailed faithful representation of reality. However, in our experience, users often want things that turn out to be incompatible:

1. They want it useful and able to satisfy their functional requirements
2. They want it the way they like it and think it should be.
3. They want it correct and theoretically sound at the same time.

Physicians are often criticized for unnecessarily using arcane and complicated language to describe clinical situations in a way inaccessible to the average patient or their family. Although physicians do sometimes purposefully obfuscate, and rarely justifiably do so, the usual reason for their choice of words derives from habit and a desire to be more precise than is possible with layman's terms. If formal ontologists are the would-be healers of terminology systems, they face an analogous criticism, and an analogous dilemma. New ideas often call for special language, but formal ontologists face a challenge at least as great as that of physicians in communicating their ideas in an accessible way to those who might make use of them. For ordinary clinicians, jargon like "endurant" and "rigid property" can mislead and obfuscate. Because of the need to have ordinary medical practitioners involved in the development and use of clinical terminology, restricting the process of clinical terminology development to a narrow group of ontological practitioners formally trained in philosophy is not a sensible way forward, and therefore the ideas of formal ontology must be communicated in a clear and understandable way. This remains an ongoing challenge.

Formal ontological principles can help to make distinctions understandable and reproducible. Some distinctions will not be useful for electronic health records or decision support, and we need to guard against a tendency towards arcane and complicated distinctions that are inaccessible to all but the most sophisticated users. However, SNOMED is actively embracing and exploiting methodologies that are shown to improve its quality and usefulness.

Acknowledgments

This work is supported in part by a grant from the College of American Pathologists.

References

[1] Welty C and Guarino N. Support for ontological analysis of taxonomic relationships. *J. Data and Knowledge Engineering*. 2001;39(1):51-74.

[2] Spackman KA, Campbell KE, Cote RA. SNOMED RT: A reference terminology for health care. *Proceedings AMIA Annual Symposium*. 1997; :640-4.

[3] College of American Pathologists. SNOMED Clinical Terms Requirements Document. 2001; http://www.snomed.org/about/documents/SNOMEDCT_Objective_V05.pdf

[4] College of American Pathologists. Commitment to Health Care. 2003;www.snomed.org/about/index.html

[5] Rector A. Clinical terminology: Why is it so hard? *Meth. Inf. Med*., 1999; 38:147-157.

[6] Campbell KE. Distributed development of a logic-based controlled medical Terminology. PhD Dissertation, Stanford University, June 1997.

[7] Spackman KA, Dionne R, Mays E, Weis J. Role grouping as an extension to the description logic of Ontylog motivated by concept modeling in SNOMED. *Proceedings AMIA Annual Symposium*. 2002; :712-716.

[8] Guarino N, Welty C, A formal ontology of properties, in R. Dieng (Ed.), *Proceedings of the 12th International Conference on Knowledge Engineering and Knowledge Management*, Springer-Verlag LNCS, Berlin, Germany, 2000

[9] Guarino N and Welty C. Evaluating Ontological Decisions with OntoClean. *Communications of the ACM*. 2002;45(2):61-65.

[10] Guarino N and Welty C. Ontological analysis of taxonomic relations. In A. Lander and V. Storey (eds.) *Proceedings of ER-2000: The International Conference on Conceptual Modeling*. 2000; Springer Verlag LNCS Vol. 1920.

[11] Fielding JM, Simon J, Smith B. Formal ontology for biomedical knowledge systems integration. http://ontology.buffalo.edu/medo/FOBKSI.pdf.

[12] Schulze-Kremer S, Smith B, Kumar A. Revising the UMLS semantic network. http://ontology.buffalo.edu/medo/UMLS_SN.pdf.

[13] Smith B, SNAP and SPAN. Actor2002 Action-oriented approaches in geographic information science. http://www.spatial.maine.edu/~actor2002/participants/smith-pres.pdf

[14] Smith B, Basic Formal Ontology and Medical Ontology, http://ontology.buffalo.edu/bfo/ BFO.htm.

[15] Campbell KE, Oliver DE, Spackman KA, and Shortliffe EH, Representing thoughts, words, and things in the UMLS. *J Am Med Inform Assoc*. 1998; 5(5): 421–431.

[16] Rector A, A tutorial on creating biomedical ontologies in DAML+OIL using OilEd http://www.cs.man.ac.uk/mig/ontology-tutorial/oiled-biomedical-ontology-tutorial.zip

[17] Schulz S, Romacker M, Hahn U. Part-whole reasoning in medical ontologies revisited - Introducing SEP triplets into classification-based description logics. *Proceedings AMIA Annual Symposium*. 1998; :830-4.