

Ontology and Medical Terminology: Why Description Logics Are Not Enough

Werner Ceusters^a, Barry Smith^b, Jim Flanagan^a

^a*Language and Computing nv., Zonnegem, Belgium*

^b*Institute for Formal Ontology and Medical Information Science, Leipzig, Germany*

Abstract

Ontology is currently perceived as the solution of first resort for all problems related to biomedical terminology, and the use of description logics is seen as a minimal requirement on adequate ontology-based systems. Contrary to common conceptions, however, description logics alone are not able to prevent incorrect representations; this is because they do not come with a theory indicating what is computed by using them, just as classical arithmetic does not tell us anything about the entities that are added or subtracted. In this paper we shall show that ontology is indeed an essential part of any solution to the problems of medical terminology – but only if it is understood in the right sort of way. Ontological engineering, we shall argue, should in every case go hand in hand with a sound ontological theory.

Keywords:

Ontology, Description Logics, Medical Terminology, Natural Language Understanding, NLP

1 Introduction

Ontology is currently perceived as the solution of first resort for all problems related to biomedical terminology, and the use of description logics is seen as a minimal requirement on adequate ontology-based systems. This is because description logics, on current conceptions, occupy the sweet-spot between maximal expressive power on the one hand and computer tractability on the other. Most ontology-based systems in healthcare accordingly conceive an ontology as a mere knowledge representation tool and thus they adhere to the computer science understanding of the term, tracing over the more venerable philosophical understanding of ontology as a science of the types of entities and relations in reality. Ontological engineering without ontological theory may, certainly, suffice for developing terminologies for data-entry systems under which users need simply to select terms while browsing through a hierarchy. But such an approach is far from meeting the requirements imposed by information systems that need to deal with grammatically complex patient records and other documents in (multiple) natural languages. In such cases, the required solution must at least contain (1) knowledge about terms and how they are used in valid constructions within natural language; (2) knowledge about the world, i.e. how the referents denoted by the terms interrelate in reality

and in given types of contexts. It must contain also (3) an algorithm that is able not only to calculate a language user's representation of that portion of the world that is described in the pertinent utterances but also (4) track the ways in which people express what does *not* represent anything in reality. For all of these purposes the required solution must be grounded in an ontological theory.

In this essay we will defend our thesis by first pointing out the problems which arise in existing medical terminological systems because they violate these principles, and then show how these principles are applied in one large-scale biomedical ontology for multilingual natural language understanding.

2 Terminological and ontological problems in medical terminologies.

Various authors have addressed the problems in current medical terminology design, focusing on aspects such as development [1] or maintenance and change [2]. Often, medical terminologies are used simply to ensure standardization of data entry. That is, pre-established hierarchies of terms are used to constrain selections made by users in annotating large document corpora. Some would like to go further and have standardized medical terminologies and associated data formats used by medical practitioners in initial recording of data. However, many clinicians share the view that faithful recording of patient data can only be achieved by using natural language. This point of view was stated already in the early 1980s by Wiederhold, who claimed that *the description of biological variability requires the flexibility of natural language and it is generally desirable not to interfere with the traditional manner of medical recording* [3]. More recently, strong arguments have been given for the use of natural language rather than controlled vocabularies in clinical records, and this from both a scientific [4] and a pragmatic [5] point of view.

Some have proposed using standardized medical terminologies for (semi-)automated natural language understanding (NLU), and it is in light of such proposals that we want to highlight some of the problems which currently existing terminologies face.

2.1 The case of SNOMED-RT

In this paper, examples are mostly drawn from the year 2000 release of SNOMED-RT (www.snomed.org). They are selected primarily in order to prove our point that the use of description logics alone is no guarantee of producing error-free systems. In [6], SNOMED RT is described as a reference terminology, i.e. "a set of concepts and relationships that provides a common reference point for comparison and aggregation of data about the entire health care process, recorded by multiple different individuals, systems, or institutions". In addition it is asserted in that document that the terminology has the following characteristics:

- It is based upon explicit representations of unique concepts, each with a single, coherent meaning, referred to by the concept's fully specified name.
- Synonyms are assigned to the ConceptId to which they refer. The exact meaning of a synonym or the exact context to which it refers can be referenced back to the concept's fully specified name.
- SNOMED RT is a multi-hierarchical terminology.
- Concepts and descriptions in SNOMED RT are represented by coded identifiers, which contain no inherent meaning or other attributes which may limit or restrict future assignment of identifiers for new terms and concepts.

- The meaning of each concept presented in SNOMED RT will never change.
- SNOMED RT supports (but does not come with) a compositional syntax that enables a single composite concept to be created by a plurality of combinations of atomic concepts and yet still be recognized as the same concept.
- SNOMED RT concepts are explicitly “defined” (we would rather say “described”, since on the basis of what SNOMED calls definitions the true meaning of a concept cannot be derived) by their semantic relationships with other concepts in the terminology. Besides the formal subsumption relationship, there are relationships such as “associated topography”, “has action”, “associated morphology” that are not themselves defined formally.

The overall structure of SNOMED-RT is shown in Figure 1.

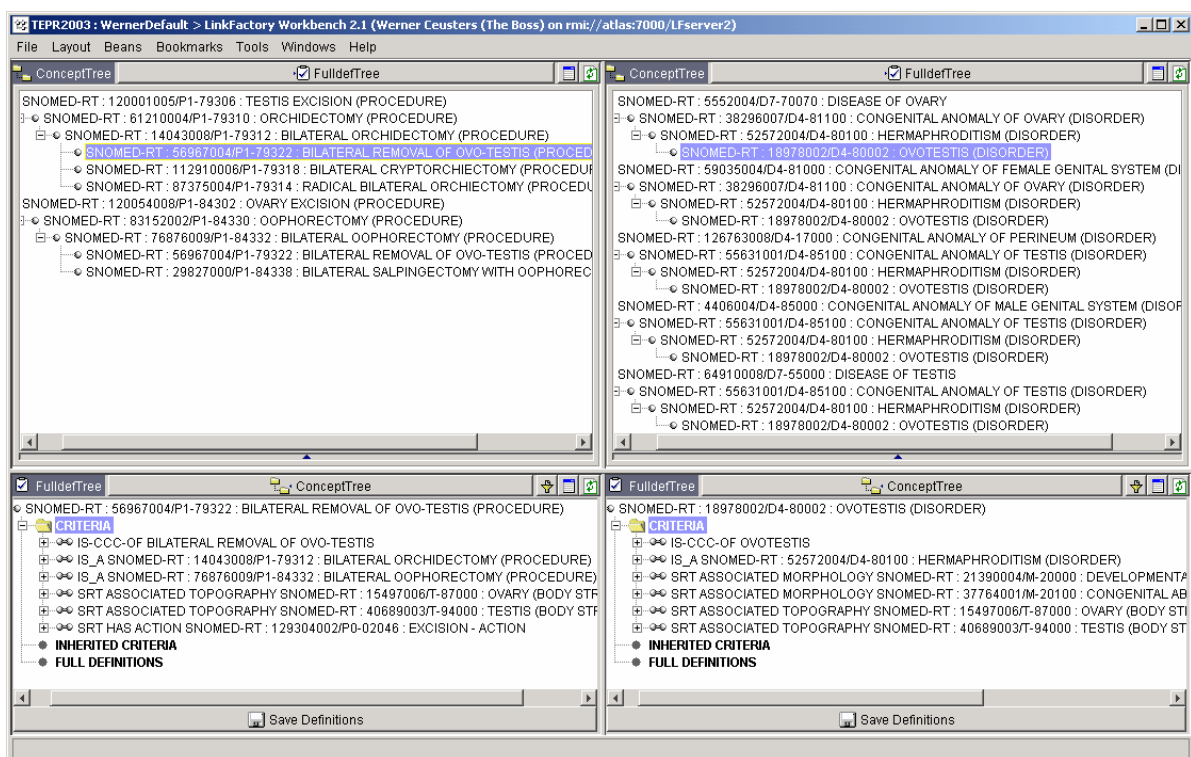


Figure 1: SNOMED-RT entries for “bilateral removal of ovo-testis” (left) and “ovo-testis” (right), showing for both entries the multiple hierarchies (top) and relationships (bottom).

Our description of problems in existing medical terminologies is organised in two parts. First, we will describe problems of a more linguistic nature. Then we will cover the problems with respect to ontology and description logics.

2.2 Linguistic problems

2.2.1 Ambiguous meanings in preferred terms

Sometimes, medical terminologies contain preferred terms that do not allow unambiguous understanding of the intended meaning when the term is viewed out of context. By simply inspecting SNOMED-RT’s hierarchy, users will likely not realize that “IMPLANTATION OF MAM-

MARY ARTERY INTO VENTRICLE (PROCEDURE)” with code 55230008/P1-31884, and “OPERATION TO ESTABLISH DRAINAGE OF VENTRICLE (PROCEDURE)” with code 34947000/P1-91050, do not refer to the same kind of “ventricle”. In the first case a cardiac ventricle is meant, in the second a cerebral ventricle. One could argue that for purposes of disambiguation the user, or in our case an NLU algorithm, should appeal to the formal concept descriptions provided by SNOMED-RT such as are shown under “criteria” in the lower part of Figure 1; but as we will show later, these very descriptions are unreliable for such purposes.

2.2.2 *Grammatical constructions in terms*

It would be nice if the totality of terms provided by a system such as SNOMED-RT could be used as a machine readable lexicon in the application of string matching algorithms to running texts. However, this is often hampered by the use in terminologies of agrammatical constructions that have their origin in old-style index construction. An example is SNOMED-RT’s 17086001/P1-48367 code that reads: “MODIFIED RADICAL MASTECTOMY, BILATERAL”. When a surgeon reports such a procedure, he would rather phrase it as “bilateral modified radical mastectomy”. This string, as it is written, does not match automatically with the original SNOMED term without further processing.

2.2.3 *Mismatch between concept label and “formal meaning”*

The terms appearing in terminologies often fail to convey what they are supposed to mean. As an example, SNOMED-RT’s concept 108034003/P1-743F9: “URINARY BLADDER EXCISION (PROCEDURE)” does not mean “excision of the urinary bladder”, but rather “excision from”, or “excision topographically related to” the urinary bladder. This can only be discovered by examining the various subsumed concepts among which one can find 53938006/P1-74340: “URINARY CYSTECTOMY (PROCEDURE)”.

Indeed, the true meaning of a SNOMED-RT concept is determined (though often only partially) by the formal relationships it bears to other concepts. This requires that before the user can understand single terms he must first understand the meanings of the various relationships employed in SNOMED-RT. This is in line with SNOMED-RT’s intended use as a reference terminology [7], but it leads to serious problems when attempts are made to use SNOMED-RT for natural language understanding without providing more detailed definitions for SNOMED concepts: for SNOMED-RT’s terms do not always correspond to what a physician means when he uses the same terms in formulating running texts.

2.3 **Ontological problems**

The world-wide attention received by the Semantic Web project as a more powerful incarnation of the current web has led to a renewed interest in “ontology”, a term whose initial currency in the information systems domain was associated primarily with the idea of ontologies as standardized classification schemes designed to advance translation between databases. It is important to note, however, that there are two importantly different uses of the term “ontology”. Most widespread is Gruber’s understanding of an ontology as a description, analogous to a formal specification of a program, of the concepts and relationships that *can exist* for an agent or a community of agents [8]. The other meaning, the one to which we adhere, sees ontology as a theory of *what exists*, a theory of the kinds and structures of objects, properties, events, processes and relations in every area of reality [9]. Ontology in this sense rests upon a number of simple principles, which include:

- objects, which *endure* through time, are categorially different from processes, which *occur* in time;
- the relation of *instantiation* (between an instance and its universal) is distinct from the relation of *subsumption* (between a less general and a more general universal);
- an *entity* is distinct from a *term* used to refer to that entity.

Quite often we come across descriptions in existing terminological systems that violate such simple ontological principles.

In SNOMED-RT, for example, the concept “testis” subsumes (correctly) the concepts “left testis”, “right testis” and “undescended testis”, but also “both testes”. We could accept *both testes* as being the mereological sum of the left and right one, but we can hardly accept it as being itself a testis. In this case, the mistake is introduced manually by one of the system builders who incorrectly assigned a direct ISA-link.

In many occasions however, erroneous subsumption is the result of computation performed by the description logic classifier used in the development process of SNOMED-RT. This is not necessarily due to errors in the algorithm, but rather to poor modelling. As an example, “ovo-testis” (a gonad that histologically contains both ovarian follicles and testicular tubular elements) is related in SNOMED-RT to “testis” and “ovary” by means of the “associated topography” relationship. Exactly the same thing is done in the concept 56967004/P1-79322: “BILATERAL REMOVAL OF OVO-TESTIS (PROCEDURE)”. However, in “ovo-testis”, SNOMED uses the associated-topography relationship more in the sense of “being there where you expect a normal ovary or a normal testis”. In “bilateral orchidectomy”, SNOMED uses the same relation to express that a testis is removed, and in “bilateral oophorectomy” that an ovary is removed (in LinKBase® we rather would use the “HAS-THEME” relationship, which is much more precise than SNOMED’s associated-topography, and which is designed to identify more explicitly the entity that moves in a motion process). As an overall result, “BILATERAL REMOVAL OF OVO-TESTIS (PROCEDURE)” becomes classified both as a “bilateral orchidectomy” and as a “bilateral oophorectomy”, since the description given for the ovo-testis removal satisfies the descriptions for both of the latter removal concepts.

Mistakes such as these are usually introduced by using relationships that are too generic. Some ontology builders, it is true, adhere to a *minimal ontological commitment* paradigm, arguing that an ontology should make as few claims as possible about the domain that is being modeled [8]. On our view, however, the job of ontology is not the construction of simplified models; rather, an ontology should correspond to reality itself in a manner that maximizes descriptive adequacy within the constraints of formal rigour and computational usefulness. Terminology authors tend to use relationships that are too generic but giving them (unconsciously) a range of more specific meanings in different types of cases. The use of “associated-topography” in SNOMED-RT, sometimes as a topological relationship and sometimes indicating participation in a motion process as described above, is a nice example. Erroneous subsumption hierarchies are one (description)-logical result. The same phenomenon can also be observed when merging several diverse terminologies into one larger system. The most prominent example is the Unified Medical Language System’s MetaThesaurus component, where a naïve peregrination up and down the hierarchical relationships leads to semantic drift and hierarchical circles [10].

We believe that many of the mistakes discussed above can be prevented by using the right sort of ontology which controls for correctness of inferences drawn by language-processing and

other tools. We first discuss the ontological theory of granular partitions [11] and then identify how this theory corresponds with the functionality provided in the ontology authoring and management system LinkFactory® [14].

3 The theory of granular partitions

The theory of granular partitions [11] is the product of an effort to build a more realistic, and also a more general and flexible, framework embodying the strengths of both set theory and mereology (the theory of part and whole), while at the same time avoiding their respective weaknesses. A *granular partition* is a grid of labeled cells. An object in the partition-theoretic sense is everything which exists that can be recognized by some cell of a partition: an individual, a part of an individual, a class of individuals (for example a biological species), a disease (as universal or as instance), a spatial region, a political unit such as a country or even the universe as a whole. Objects can be either of the bona fide or of the fiat sort [12]. Bona fide objects, for example a patient's body, are objects which exist and are demarcated from their surroundings independently of human partitioning activity. Fiat objects are objects which exist and are demarcated from their surroundings only because of such partitioning activity. An example would be your right arm: where it starts from the shoulder is a matter of drawing a demarcation line that does not exist in reality. Other examples are: the borders between the midbrain, pons, medulla and spinal cord. Partitions are more than just systems of cells. They are constructed to serve as inventories or maps of specific portions of reality, and in this they are analogous to *windows*, or to the latticed grills purported to have been used by Renaissance artists as aids to the faithful representation of objects in reality [13].

The main principles of the theory can be summarized as follows:

- a partition is the drawing of a (typically complex) fiat boundary over a certain domain
- a partition typically comes with labels and/or an address system
- partitions are artifacts of our cognition
- a partition is transparent (it should be veridical, should neither simplify nor change reality)
- bona fide objects exist independently of our partitions, fiat objects are determined by partitions
- different partitions may represent cuts through the same reality which are skew to each other
- entities (existing in reality) located in the same cell of a partition share common characteristics
- partitions can allow for the addition of new information and new cellular divisions reflecting growth in knowledge about the relevant domain

Partonomies, taxonomies and maps are all specific examples of granular partitions.

4 Ontological engineering in LinKBase®

LinKBase® is a large scale medical ontology developed and maintained by the modeling team of Language and Computing nv and using the ontology authoring system LinkFactory®. LinKBase® contains currently over 1.5 million language-independent medical and general-

purpose universals (such as “human body”, “molecule”, “surgical dilatation”) as well as particular instances (such as “United States”, “WHO”), all linked to natural language terms in several languages [14, 15]. These designations for universals and particulars are linked together into a semantic network-like structure using approximately 450 different link-types for expressing formal relationships in a framework which draws on logically axiomatized theories in mereology and topology [16, 17] as well as theories of such notions as time and causality [18]. In addition, the framework has tools for semantics-driven natural language understanding [19, 20].

Link-types form a multi-parented hierarchy in their own right. It is important to note that in LinKBase® the formal subsumption relationship covers only about 15% of the total number of relationships. As such, LinKBase® is a much richer structure than terminological systems in which term-relationships are expressed as subsumptions (i.e. as relations between “narrower” and “broader” terms). It is important to note also that LinKBase® is a living ontology, in which data are changed on a daily basis and at a rate of some 2000 to 4000 modifications a day. Moreover, it is not required for universals added to be perfectly described from the very beginning [21].

4.1 Partitions in LinKBase®




Figure 2. Three partitions of the domain meningitis.

LinKBase® describes reality using four central representational devices. The first is: “domain-entity”, which subsumes any entity in reality that has been classified by the ontology authors. As such, the notion of “domain-entity” corresponds within the theory of granular partitions to the notion of an entity that is recognized by the cell of a partition. Examples are MENINGITIS (as captured by the inflammation partition) and MENINGES (as captured by an anatomical map).

The term “meta-entity” refers to entities or putative entities as these have been foregrounded by cells of partitions prepared by third parties. An example is the (meta-)entity referred to by the ICD-10 term “G03.9”, standing for MENINGITIS, UNSPECIFIED. Groups of interrelated meta-entities in LinKBase® constitute views on reality adhered to by external authors. For these partitions, it is assumed neither that they are veridical nor that they are not veridical.

The third representational device is a domain-entity-link, which is used to produce descriptions relating one domain-entity to other domain-entities. An example is HAS-SYSTEMIC-MEDIUM, which relates a domain-entity of the type *process* to the domain-entity that is a central participant in a process, such as the heart in cardiac hypertrophy, or the meninges in meningitis. Domain-entity-links represent a specific aspect of domain-entities determining a partition or perspective. Terms of domain-entity-links can be both universals and instances.

The last central notion is that of a criterion, which can be thought of as a combination of a domain-entity-link with the domain-entity towards which the link points. In the theory of granular partitions, this corresponds to the perspective on a domain-entity that is yielded by a given partition. Examples are HAS-SYSTEMIC-MEDIUM MENINGES and HAS-CONSEQUENCE HEADACHE.

4.2 How statements in LinkBase® correspond to entities in reality

A detailed description of the ontological conventions to which LinkFactory® adheres is given in [14], [15], and [21]. Only those that can easily be expressed in terms of granular partition theory are given here.

Crucial for a correct use of LinkFactory® is the requirement that, if a real-world entity is declared to be an instance of a domain-entity, then all that is said about the domain-entity necessarily applies to the instance (indeed to all instances of the same universal). In line with the principle that partitions (other than those included under the heading “Meta-Entities”) are *veridical* – that they should neither create nor change reality – this means that the real-world-entity is required to be located in the correct cells of all partitions corresponding to the domain-entity-links used in a domain-entity’s description. It is important to understand that the ontology authors are responsible for creating veridical partitions and that LinkFactory® is designed to help them in realizing this goal.

LinkFactory® can represent necessary and sufficient conditions for a real-world-entity to be identified as an instance of a domain-entity, and this on the basis of information that has been collected about the real-world-entity. Note that identification is distinct from declaration as described in the previous paragraph. While declaration allows you to assign everything that is valid for a universal *U* to be valid for a particular *p* if you know that *p IS-INSTANCE-OF U*, identification allows you to infer that *p IS-INSTANCE-OF U*, when you know at least that *p* satisfies the necessary and sufficient conditions associated with *U*. As such, LinkFactory® permits partitions to be used as knowledge elicitation devices.

LinkBase®’s domain-entity-links have a direction: one can say that they point from a source concept to a target concept. In a LinkBase® statement such as “MENINGITIS HAS-SYSTEMIC-MEDIUM MENINGES” the source concept is “meningitis”, and the target concept is “meninges”. It is crucial to understand that in the formalism underlying LinkFactory®, this statement does not tell us anything about meninges unless the domain-entity-link “has-systemic-medium” is declared to have an inverse link.

Inverse links come in two flavours. If a link is declared “autocontra” with an inverse link such as in

“IS-UNDER” – autocontra – “IS-ABOVE”

then any statement A IS-UNDER B (meaning that each instance of the universal A is under some instance of the universal B), entails B IS-ABOVE A (meaning each instance of the universal B is above some instance of the universal A).

If however a link is declared “contra” with an inverse link such as in

“HAS-CAUSE” – contra – “IS-CAUSE-OF”

then a statement such as “bacterial meningitis” HAS-CAUSE “bacterium” does not entail that for all instances of “bacterium” there is some meningitis of which they are the cause. Only for a specific instance of “bacterial meningitis”, that obviously HAS-CAUSE a specific bacterium, does it hold that some specific bacterium IS-CAUSE-OF that specific bacterial meningitis. This corresponds to the notion that there are partitions that allow you to see organisms such as bacteria without necessarily allowing you to see also the associated disorders. The combined use of partitions for body parts and organisms ought under certain conditions to allow the

construction of partitions enabling you to see associated diseases. This is an issue that must be further explored [11].

5 Conclusion

Too many (medical) ontologies have been designed without keeping linguistic or ontological constraints in mind. That systems designed without any true ontological background suffer from problems hampering their use in advanced natural language understanding applications does not come as a surprise. An example is UMLS and its components, as shown by statements such as “*Simply using everything in the Metathesaurus does not make a good coding system*” [22], and:

The problems with the Metathesaurus as a single monolithic vocabulary are: 1. There is a wide range of granularity of terms in different vocabularies. 2. The Metathesaurus itself has no unifying hierarchy, so you cannot take advantage of hierarchical relations. 3. There may be other features of vocabularies that get lost in their ‘homogenisation’ upon being entered into the Metathesaurus. [23].

Less obvious for people who are not experts in ontology or description logics is the fact that “doing (some sort of) ontology” or “using description logics” is not a guarantee for accurate system development. As Elkin *et al.* expressed it:

The current implementation of SNOMED-RT does not have the depth of semantics necessary to arrive at comparable data or to algorithmically map to classifications such as ICD-9-CM” [24, p. 70].

There are different sorts of specializations employed in medical documentation. Thus while a concept such as “removal of liver” has a primary meaning of the entire liver being removed, it admits of specializations such as “removal of liver lobe”. The possibility of such different sorts of specialization needs to be built into the ontology one uses in ways which go beyond conceptual subsumption, which is the sort of specialization admitted by description logics. As Hahn *et al.* point out: “*A serious limitation of the Galen approach is that specialisation is invariably linked to a conceptual relation*” [25, p. 66].

The only good approach is to have systems that keep natural language and ontology nicely in balance. It is not sufficient to have formal tools that do not rely on a sound ontological theory. If clinicians can understand Francis Bacon’s saying “*He who will not apply new remedies, must expect old evils*”, why should not the developers of medical ontologies do so also?

Acknowledgement: Work on this paper was supported by the Alexander von Humboldt Foundation under the auspices of its Wolfgang Paul Programme.

References

- [1] **Cimino JJ.** *Desiderata for controlled medical vocabularies in the twenty-first century.* Methods of Information in Medicine, 37(4-5):394-403, 1998.
- [2] **Oliver DE., Shahar Y., Musen MA., & Shortliffe EH.** *Representation of Change in Controlled Medical Terminologies.* Artificial Intelligence in Medicine 15(1):53-76, 1999.
- [3] **Wiederhold G.** *Databases in healthcare.* Stanford University, Computer Science Department, Report No. STAN-CS-80-790, 1980.
- [4] **Kay S, Purves IN.** *Medical Records and Other Stories: a narratological framework.* Methods of Information in Medicine 1996; 35: 72-87.

- [5] **Blum E.** *Paperless medical record not all it's cracked up to be.* Commentary in AMNews. Feb. 17, 2003 (http://www.ama-assn.org/sci-pubs/amnews/pick_03/bica0217.htm).
- [6] **College of American Pathologists.** *SNOMED RT - Systematized Nomenclature of Medicine Reference Terminology, VERSION 1.1, USER GUIDE*, 2001.
- [7] **Spackman KA, Campbell KE, Cote RA.** *SNOMED RT: A reference terminology for health care.* In DR Masys (Ed.), JAMIA 1997, Proceedings Supplement: AMIA Annual Fall Symposium, pp. 640-644.
- [8] **Gruber T.R..** *A translation approach to portable ontologies.* Knowledge Acquisition, 5(2):199-220, 1993
- [9] **Smith B and Welty C.** *Ontology: Towards a new synthesis.* In Chris Welty and Barry Smith, eds., *Formal Ontology in Information Systems*. Pp. iii-x. Ogunquit, Maine: ACM Press. 2001.
- [10] **Bodenreider O.** *Circular Hierarchical Relationships in the UMLS: Etiology, Diagnosis, Treatment, Complications and Prevention.* Proceedings of AMIA Annual Symposium 2001:57-61.
- [11] **Bittner Th and Smith B.** *A Theory of Granular Partitions.* In: Foundations of Geographic Information Science, Matthew Duckham, Michael F. Goodchild and Michael F. Worboys, eds., London: Taylor & Francis, 2003, 117–151.
- [12] **Smith B.** *Fiat objects.* Topoi, 20(2):131–148, 2001.
- [13] **Smith B.** *True grid.* In Montello, D. R. (Ed), *Spatial Information Theory: Foundations of Geographic Information Science*, volume 2205 of Lecture Notes in Computer Science, pp. 14–27. Berlin: Springer-Verlag.
- [14] **Ceusters W, Martens P, Dhaen C, Terzic B,** *LinkFactory: an Advanced Formal Ontology Management System.* Interactive Tools for Knowledge Capture Workshop, KCAP-2001, October 20, 2001, Victoria B.C., Canada (<http://sern.ucalgary.ca/ksi/K-CAP/K-CAP2001/>).
- [15] **Montyne F,** *The importance of formal ontologies: a case study in occupational health.* OES-SEO2001 International Workshop on Open Enterprise Solutions: Systems, Experiences, and Organizations, Rome, 14-15 September 2001 (<http://cersi.luiss.it/oeseo2001/papers/28.pdf>).
- [16] **Smith B,** *Mereotopology: a theory of parts and boundaries,* Data and Knowledge Engineering 20 (1996), 287-301.
- [17] **Smith B, Varzi AC,** *Fiat and Bona Fide Boundaries,* in Proc. COSIT-97, Springer-Verlag 1997, 103-119. Revised version in *Philosophy and Phenomenological Research*, 60: 2, March 2000, 401–420.
- [18] **Buekens F, Ceusters W, De Moor G,** *The Explanatory Role of Events in Causal and Temporal Reasoning in Medicine,* Met Inform Med 1993, 32: 274 - 278.
- [19] **Ceusters W, Buekens F, De Moor G, Waagmeester A,** *The distinction between linguistic and conceptual semantics in medical terminology and its implications for NLP-based knowledge acquisition.* Met Inform Med 1998; 37(4/5):327-33.
- [20] **Bateman JA.** *Ontology construction and natural language.* In Proc. International Workshop on Formal Ontology. Padua, Italy, 1993, 83-93.
- [21] **Flett A, Casella dos Santos M, Ceusters W.** *Some Ontology Engineering Processes and their Supporting Technologies,* in: Gomez-Perez A, Benjamins VR (eds.) *Ontologies and the Semantic Web*, EKAW2002, Springer 2002, 154-165.
- [22] **William T. Hole** M.D., Director, Metathesaurus Research and Development, National Library of Medicine. Message to UMLS-users Mailing List, 23-06-2000.
- [23] **Hersh W.** Message to UMLS-users Mailing List, 23-06-2000.
- [24] **Elkin PL, Harris M, Ogren PV, Buntrock ID, Brown SH, Solbrig HR, Chute CG:** *Semantic Augmentation of Description Logic based Terminologies* Addendum to Proceedings of IMIA-WG6, Medical Concept and Language Representation, Phoenix, 16-19/12/1999, 70-81.
- [25] **Hahn U, Schulz S, Romacker M:** *Part-whole reasoning: a case study in medical ontology engineering.* IEEE Intelligent Systems & Their Applications vol 14 nr 5, 1999: 59-67.

Dr. Werner Ceusters & Dr. Jim Flanagan, Language & Computing nv, Hazenakkerstraat 20a, B9520 – Zonnegem, Belgium, www.landcglobal.com

Prof. Barry Smith, Institute for Formal Ontology and Medical Information Science, University of Leipzig, Härtelstraße 16-18, 04107 Leipzig, Germany, <http://www.ifomis.de> and Department of Philosophy, University at Buffalo, NY, USA.,