Towards a Broad-Coverage Biomedical Ontology Based on Description Logics

U. Hahn, S. Schulz

Pacific Symposium on Biocomputing 8:577-588(2003)

TOWARDS A BROAD-COVERAGE BIOMEDICAL ONTOLOGY BASED ON DESCRIPTION LOGICS

U. HAHN [a] S. SCHULZ [a,b]

[a] *Text Knowledge Engineering Lab, Freiburg University, D-79085 Freiburg, Germany*
`http://www.coling.uni-freiburg.de`
[b] *Department of Medical Informatics, University Hospital, D-79104 Freiburg, Germany*

We describe an ontology engineering methodology by which conceptual knowledge is extracted from an informal medical thesaurus (UMLS) and automatically converted into a formal description logics system (LOOM). Our approach consists of four steps: concept definitions are automatically generated from the UMLS, integrity checking of taxonomic and partonomic hierarchies is performed by LOOM's terminological classifier, cycles and inconsistencies are eliminated, as well as incremental refinement of the evolving knowledge base is performed by a domain expert. We report on experiments with a very large knowledge base composed of 164,000 concepts and 76,000 relations.

1 Introduction

Unlike many other disciplines, medicine has a long standing tradition in structuring its domain knowledge, e.g, disease taxonomies, medical procedures, anatomical terms, in a wide variety of medical terminologies, thesauri and classification systems. These efforts are typically restricted to the provision of broader and narrower terms, related terms or synonymous terms. This is most evident in the UMLS, the *Unified Medical Language System*[1], an umbrella system which covers more than 60 medical thesauri and classifications (e.g., MeSH, ICD, SNOMED, Digital Anatomist).

From a conceptual perspective, the UMLS can be divided into two major parts. The UMLS *Semantic Network (SN)*, on the one hand, forms the upper ontology and consists of 134 semantic types linked by 54 types of semantic relations, which makes a total of 7,473 edges, The UMLS *Metathesaurus*, on the other hand, contains 776,940 concepts (2002 version), each of which is assigned to one or more UMLS SN types. These concepts are linked by the semantic relations also supplied by the UMLS SN. In total, 10,147,419 semantic links between these Metathesaurus concepts exist, most of them directly taken from the sources, some added by the UMLS developers. The vast majority of these links introduce thesaurus-like broader/narrower term relationships.

Both, the UMLS SN and the Metathesaurus, form a huge semantic network. Its semantics is shallow and entirely intuitive, which is due to the fact that their usage was primarily intended for humans in order to support health-related knowledge management. Given the size, the evolutionary diversity and inherent heterogeneity of the UMLS, there is no surprise that the lack of a formal semantic foundation leads to inconsistencies, circular definitions, etc.[2,3]. This may not cause utterly severe problems when humans are in the loop and its use is limited to disease or procedure encoding, accountancy or document retrieval tasks. However, anticipating its use for more knowledge-intensive applications such as medical decision making or the understanding of medical narratives those shortcomings might lead to an impasse.

As a consequence, formal models for dealing with medical knowledge have been proposed such as conceptual graphs, semantic networks or description logics[4,5,6,7]. Not surprisingly, there is a price to be paid for more expressiveness and formal rigor, *viz.* increasing modeling efforts and, hence, increasing maintenance costs[8]. Operational systems making full use of such rigid approaches, especially those which employ high-end knowledge representation languages, are usually restricted to rather small subdomains. The most comprehensive of these sources we know of is the GRAIL-encoded GALEN knowledge base which covers up to 9,800 concepts[6]. The limited coverage then hampers their routine usage, an issue which is always highly rewarded in the medical informatics community.

Almost all of the knowledge bases developed on the basis of formal representation languages have been designed from scratch – without making systematic use of the large body of knowledge contained in widely spread medical terminologies. Hence, it would be an intriguing approach to join the massive *coverage* offered by informal medical terminologies with the high level of *expressiveness* and *reasoning capabilities* supported by rigid knowledge representation systems in order to develop formally solid medical knowledge bases on a larger scale. This idea has already been fostered by Pisanelli et al.[9] who extracted knowledge from the UMLS SN and from parts of the Metathesaurus, and merged them with logic-based top-level ontologies from various sources. Another example is the re-engineering of SNOMED[10] from a multi-axial coding system into a formally founded ontology[11,12].

Unfortunately, the efforts made so far are entirely focused on generalization-based reasoning along taxonomies and lack a reasonable coverage of partonomies. As the latter form a crucial part of medical knowledge, modeling efforts not only have to be directed at generalization hierarchies and taxonomic reasoning but must incorporate part-whole knowledge and partonomic reasoning, as well. We here describe such an integrated knowledge engineering methodology. The resulting medical ontology will form the domain knowledge backbone of MEDSYNDIKATE, a system for the automatic acquisition of knowledge from medical finding reports[13].

2 Reasoning Along Part-Whole Hierarchies

Medical knowledge is typically organized around generalization hierarchies – on which taxonomic reasoning along *is-a* relations is based – and part-whole hierarchies – allowing partonomic reasoning along *part-of* or *has-part* relations. Unlike generalization-based reasoning in concept taxonomies, no fully conclusive mechanism exists up to now for reasoning along part-whole hierarchies (for a survey of different approaches, cf. Artale et al.[14]).

Within the description logics (DL) paradigm several extensions to representation languages have been proposed which provide special constructors for part-whole reasoning[6,15]. This seems a reasonable way to proceed as long as the transitivity property of a relation can be assumed, in general. In the medical[16], as well as in commonsense domains[17,18], however, this view has been invalidated. Obviously, various exceptions exist such that the transitivity of *part-of* relations cannot be granted, in general. Hence, both the expression of regular transitive use, as well as exception handling for nontransitive *part-of* relations have to be taken into consideration. Another challenge comes with the propagation of properties across part-whole hierarchies, often referred to as 'inheritance across transitive roles' (e.g., *inflammation-of* ∘ *part-of* → *inflammation-of*)[19]. Especially with biomedical knowledge this reasoning pattern cannot be generalized, since it faces lots of exceptions, too.

Motivated by previous approaches[20,21], we formalized a model of partonomic reasoning[16] that accounts for the above considerations and, also, does not exceed the expressiveness of the well-understood, parsimonious concept language $\mathcal{ALC}$[22].[a] Our proposal is centered around a particular data structure for partonomic reasoning, so-called *SEP triplets* (cf. Figure 1). They define a characteristic pattern of *is-a* hierarchies which support the emulation of inferences typical of transitive *part-of* relations. In this formalism, the relation *anatomical-part-of* describes the partitive relation between physical parts of an organism.

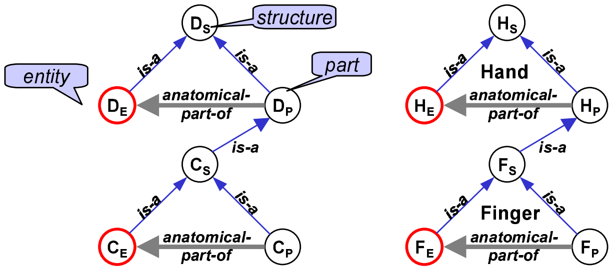

Figure 1: SEP Triplets: Partitive Relations within Taxonomies

A triplet consists, first of all, of a composite *Structure* concept, the so-called *S-node* (e.g., *Hand-Structure*, H_S). Each *Structure* concept subsumes directly an anatomical *Entity* concept, on the one hand, and a common subsumer of anything that

---

[a] $\mathcal{ALC}$ allows for the construction of concept hierarchies, where '$\sqsubseteq$' denotes subsumption and '$\doteq$' definitional equivalence. Existential ($\exists$) and universal ($\forall$) quantification, negation ($\neg$), disjunction ($\sqcap$) and conjunction ($\sqcup$) are also supported. Role filler constraints (e.g., typing by C) are linked to the relation name R by a dot, $\exists R.C$.

is a *Part* of that entity concept, on the other hand. These two concepts are called *E-node* and *P-node*, e.g., *Hand-Entity* (H_E) and *Hand-Part* (H_P), respectively. Whereas E-nodes denote the anatomical concepts proper to be modelled in our domain, S-nodes and P-nodes constitute representational artifacts required for the formal reconstruction of the systematic patterns and exceptions underlying partonomic reasoning. More precisely, a P-node is the common subsumer of those concepts that have their role *anatomical-part-of* filled by the corresponding E-node concept, as an existential condition. For example, *Hand-Part* subsumes those concepts all instances of which have a *Hand-Entity* as a necessary whole. As an additional constraint, E-nodes and P-nodes can be modelled as being mutually disjoint. This is a reasonable assumption for most concepts denoting singleton objects, where parts and wholes cannot be of the same type (a red blood cell cannot be part of yet another red blood cell). On the contrary, masses and collections can have parts and wholes of the same type, e.g., a tissue can be part of another tissue.

For the reconstruction of the *anatomical-part-of* relation by taxonomic reasoning, we assume C_E and D_E to denote E-nodes, C_S and D_S to denote the S-nodes that subsume C_E and D_E, respectively, and C_P and D_P to denote the P-nodes related to C_E and D_E, respectively, via the role *anatomical-part-of* (cf. Figure 1). These conventions can be captured by the following terminological expressions:

$$C_E \sqsubseteq C_S \sqsubseteq D_P \sqsubseteq D_S \tag{1}$$

$$D_E \sqsubseteq D_S \tag{2}$$

The P-node is defined as follows (we here introduce the disjointness constraint between D_E and D_P, i.e., no instance of D can be *anatomical-part-of* any other instance of D):

$$D_P \doteq D_S \sqcap \neg D_E \sqcap \exists anatomical\text{-}part\text{-}of.D_E \tag{3}$$

Since C_E is subsumed by D_P (according to (1)), we infer that the relation *anatomical-part-of* holds between C_E and D_E, too:[b]

$$C_E \sqsubseteq \exists anatomical\text{-}part\text{-}of.D_E \tag{4}$$

The encoding of concept hierarchies in terms of SEP triplets allows the knowledge engineer to switch the transitivity property of part-whole relations off and on, dependent on whether the E-node or the S-node, respectively, is addressed as the target concept for a conceptual relation. In the first case, the propagation of roles

---

[b] An extension of this encoding scheme which allows additional reasoning about *has-part* in a similar way, is proposed in[23], though it has not been considered in the knowledge base described in this paper.

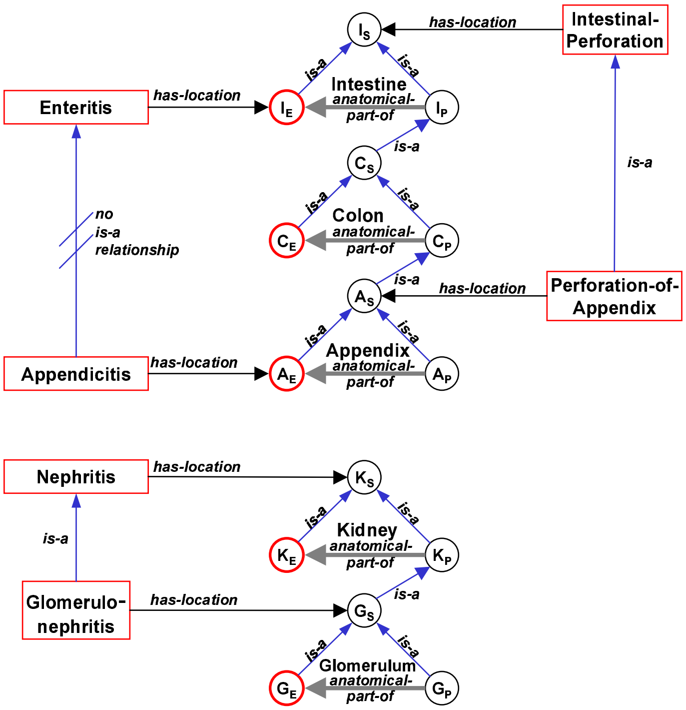

Figure 2: Enabling/Disabling Role Propagation in a SEP-Encoded Partonomy

across part-whole hierarchies is disabled, in the second case it is enabled. As an example (cf. Figure 2), *Enteritis* is defined as *has-location Intestine*$_E$, i.e., the range of the relation *has-location* is restricted to the E-node of *Intestine*. This precludes, e.g., the classification of *Appendicitis* as *Enteritis* though the *Appendix* is related to the *Intestine* via an *anatomical-part-of* relation. In the 'switch-on' mode, however, *Glomerulonephritis* (*has-location Glomerulum*$_S$) is classified as *Nephritis* (*has-location Kidney*$_S$), with *Glomerulum* being an *anatomical-part-of* the *Kidney*. In the same way, *Perforation-of-Appendix* is classified as *Intestinal-Perforation* (cf. Hahn et al.[16] for an in-depth analysis of these phenomena).

3 Ontology Engineering Workflow

Our goal is to extract conceptual knowledge from two major subdomains of the UMLS, *viz.* anatomy and pathology, in order to construct a formally sound knowledge base based on $\mathcal{ALC}$-type description logics[22]. Such a research program puts considerable pressure on well-engineered terminological reasoning systems, since we require their inference engine to perform classification-based reasoning reliably on very large data sets ($\gg$ 100,000 items). Hence, we consider the LOOM system[24], at the system level, the most convenient match of our requirements and the state of the art in terminological reasoning (for an empirical study, cf. Heinsohn et al.[25]).

The ontology engineering task will be divided into four steps: (1) automatic generation of terminological expressions, (2) automatic consistency checking by a terminological classifier, (3) manual restitution of formal consistency in case of inconsistencies, and, (4) manual curation of the formal representation structures.

Step 1: Automatic Generation of Terminological Expressions. Sources for concepts and relations were the 1999 release of the UMLS SN and the UMLS metathesaurus. Figure 3 exhibits the semantic links between two UMLS CUIs (concept unique identifier).[c] These tables, available as ASCII files, were imported into a Microsoft Access relational database and manipulated using SQL embedded in the VBA programming language. For each CUI in the *mrrel* subset its alphanumeric code was substituted by the English preferred term.

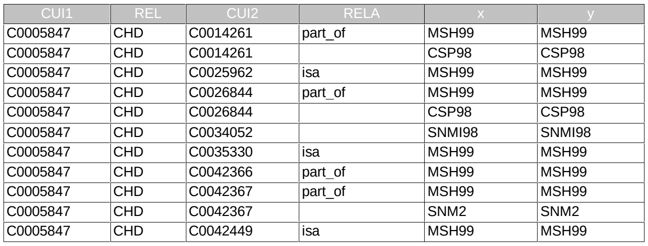

CUI1	REL	CUI2	RELA	x	y
C0005847	CHD	C0014261	part_of	MSH99	MSH99
C0005847	CHD	C0014261		CSP98	CSP98
C0005847	CHD	C0025962	isa	MSH99	MSH99
C0005847	CHD	C0026844	part_of	MSH99	MSH99
C0005847	CHD	C0026844		CSP98	CSP98
C0005847	CHD	C0034052		SNMI98	SNMI98
C0005847	CHD	C0035330	isa	MSH99	MSH99
C0005847	CHD	C0042366	part_of	MSH99	MSH99
C0005847	CHD	C0042367	part_of	MSH99	MSH99
C0005847	CHD	C0042367		SNM2	SNM2
C0005847	CHD	C0042449	isa	MSH99	MSH99

Figure 3: Semantic Relations in the UMLS Metathesaurus

After manual remodeling of the top-level concepts of the UMLS SN we extracted, from a total of 85,899 metathesaurus concepts, 38,059 anatomy and 50,087 pathology concepts. The criterion for the inclusion into one of these sets is the assignment to predefined UMLS SN types. Also, 2,247 concepts appeared in both sets, anatomy and pathology. Since we wanted to keep the two subdomains strictly disjoint, we duplicated these overlapping concepts, and prefixed all concepts by *ana*- or *pat*- according to their proper subdomain. Indeed, these hybrid concepts mirror multiple meanings. For instance, *tumor* has the meaning of a malignant disease, on the one hand, and of an anatomical structure, on the other hand.

---

[c] As a convention in UMLS, any two CUIs must be connected by at least a shallow relation (in Figure 3, CHilD relations in the column REL are assumed between CUIs). Shallow relations may be refined in the column RELA, if a thesaurus is available which contains more specific information. Some CUIs are linked either by *part-of* or *is-a*. In any case, the source thesaurus for the relations and the CUIs involved is specified in the columns X and Y (e.g., MeSH 1999, SNOMED International 1998).

As target structures for the anatomy domain we chose SEP triplets. Only UMLS-supplied *part-of* (including *conceptual-part-of*), *has-part* (including *has-conceptual-part*) and *is-a* relation attributes were considered for the construction of taxonomic and partonomic hierarchies. Hence, for each anatomy concept one SEP triplet was created. For the pathology domain, we treated *CHD* (*child*) and *RN* (*narrower* relation) from the UMLS as indicating taxonomic (*is-a*) links. No part-whole relations were considered, since there were only two occurences of the relation attribute *part-of* between pathology concepts in the whole metathesaurus. Furthermore, for all anatomy concepts contained in the definitional statements of pathology concepts the S-node is the default concept to which they are linked, thus enabling the propagation of roles across the part-whole hierarchy.

As a fundamental assumption, all roles generated in this process were considered as being existentially quantified. This means that any relation r (*part-of*, *has-location*, etc.) which holds between two concepts, A and B, is mapped to a corresponding role $\tilde{r}.B$ which is a necessary condition in the definition of the concept A. All conceptual constraints for a concept definition are mapped to a conjunction of constraints.

Another basic assumption concerns the strict antisymmetry of *is-a* and *part-of*. As a consequence, we did not process reflexive relations such as "C1 *broader-than* C1" (n = 9,546 cases of reflexive hierarchical relations were found).

Furthermore, type-to-type relations from the UMLS SN (e.g., *Anatomical Structure location-of Disease or Syndrome*) were not considered, since they indicate "possible" relations that cannot be adequately expressed in description logics.

Step 2: Automatic Consistency Checking by the LOOM **Classifier.** The import of UMLS anatomy concepts resulted in 38,059 *deftriplet* expressions for anatomical concepts and 50,087 *defconcept* expressions for pathological concepts. Each *deftriplet* was expanded into three *defconcept* (S-, E-, and P-nodes), and two *defrelation* (*anatomical-part-of-x*, *inv-anatomical-part-of-x*) expressions, summing up to 114,177 concepts. This yielded (together with the concepts from the UMLS SN) a total of 240,764 definitory LOOM expressions.

From 38,059 anatomy triplets, 1,219 *deftriplet* statements contained a *:has-part* clause followed by a list of a variable number of triplets, with more than one argument in 823 cases (3.3, on average). 4,043 *deftriplet* statements contained a *:part-of* clause, only in 332 cases followed by more than one argument (1.1, on average). We then submitted the resulting knowledge base to the terminological classifier, the inference engine which computes subsumption relations, and checked for terminological cycles and consistency. In the anatomy subdomain, one terminological cycle and 2,328 inconsistent concepts were found, in the pathology subdomain 355 terminological cycles though not a single inconsistent concept were determined.

Step 3: Manual Restitution of Consistency. The inconsistencies in the anatomy part of the knowledge base identified by the classifier could all be traced back to the simultaneous linkage of two triplets by both *is-a* and *part-of* links, an encoding that raises a conflict due to the disjointness required for corresponding P- and E-nodes (cf. expression (3)). In most of these cases the parents involved belonged to a class of concepts that obviously cannot be appropriately modeled as SEP triplets, e.g., *Subdivision-Of-Ascending-Aorta* or *Organ-Part*. The meaning of each of these concepts almost paraphrases that of a P-node, so that the violation of the SEP-internal disjointness condition could be accounted for by substituting the triplets involved with simple LOOM concepts, by matching them with already existing P-nodes, or by disabling *is-a* or *part-of* links. B In the pathology part of the knowledge base, we expected a large number of terminological cycles to occur, simply as a consequence of interpreting the extremely weak *narrower term* and *child* relations in terms of taxonomic subsumption (*is-a*). Bearing in mind the size of the knowledge base, we consider 355 cycles quite a low number (Bodenreider[3] found a total of 1,920 cycles in the complete 2001 metathesaurus release.). Those cycles were primarily due to very similar concepts, e.g., *Arteriosclerosis vs. Atherosclerosis*, *Amaurosis vs. Blindness*, and residual categories ("other", "NOS" = *not otherwise specified*). These concepts were directly inherited from the source terminologies and are notoriously difficult to interpret out of their definitional context, e.g., *Other-Malignant-Neoplasm-of-Skin vs. Malignant-Neoplasm-of-Skin-NOS*. In many cases the decision which relations could be maintained and which ones had to be eliminated was taken arbitrarily, since for biomedical terminology often no consensus can be reached on the exact meaning of terms. As the result of the analysis we obtained a negative list which consisted of 630 concept pairs. In a subsequent extraction cycle we incorporated this list in the automated construction of the LOOM concept definitions and, with these new constraints added, a fully consistent knowledge base was generated, finally.

Step 4: Manual Curation of the Knowledge Base. To set up this high-volume knowledge base including the aforementioned working steps required three months of work for a single person. The fourth step – when performed for the whole knowledge base – is very time-consuming and requires broad and in-depth medical expertise. Random samples from both subdomains were analyzed by the second author, a domain expert. The data we here supply refer to the analysis of two random samples of each 100 anatomy and 100 pathology concepts. This took one person about a single month. From the experience we gained so far, the following workflow can be derived:

- *Checking the correctness of the taxonomic and partonomic hierarchies.* Taxonomic and partonomic links are manually added or removed.

  Results: In the *anatomy* sample, only 76 in 100 concepts could be unequivocally classified as belonging to 'canonical' anatomy. (The remainder, e.g., *ana-Phalanx-of-Supernumerary-Digit-of-Hand*, referring to pathological anatomy, was immediately excluded from analysis.) Besides the assignment to the UMLS semantic types, only 27 (direct) taxonomic links were found. 83 UMLS relations (mostly *child* or *narrower* relations) were manually upgraded to taxonomic links. 12 (direct) *part-of* and 19 *has-part* relations were found. Four *part-of* relations and one *has-part* relation had to be removed, since we considered them as implausible. 51 UMLS relations (mostly *child* or *narrower* relations) were manually upgraded to *part-of* relations, and 94 UMLS relations (mostly *parent* or *broader* relations) were upgraded to *has-part* relations. After this workup and upgrade of shallow UMLS relations to semantically more specific relations, the sample was checked for completeness again. As a result, 14 *is-a* and 37 *part-of* relations were still considered as missing.

In the *pathology* sample, the assignment to the pathology subdomain was considered plausible for 99 of 100 concepts. A total of 15 false *is-a* relations was identified in 12 concept definitions. 24 *is-a* relations were found to be missing.

- *Check of the* :has-part *arguments assuming 'real anatomy'.* In the UMLS sources *part-of* and *has-part* relations are considered as symmetric. According to our transformation rules, the attachment of a role *has-anatomical-part* to an E-node B_E, with its range restricted to A_E, implies the existence of a concept *A* for the definition of a concept *B*. On the other hand, the classification of A_E as subsumee of the P-node B_P, the latter being defined via the role *anatomical-part-of* restricted to B_E, implies the existence of B_E given the existence of A_E.

  These constraints do not always conform to 'real' anatomy, i.e., anatomical entities that may exhibit pathological modifications. Figure 4 (left) sketches a concept A that is necessarily *anatomical-part-of* a concept B, but whose existence is not required for the definition of B. This is typical of the results of surgical interventions, e.g., a large intestine without an appendix, or an oral cavity without teeth, etc.

  Results: All 112 *has-part* relations obtained by the automatic import and the manual workup of our sample were checked. The analysis revealed that more than half of them (62) should be eliminated in order not to obviate a coherent classification of pathologically modified anatomical objects. For instance, maintaining *has-anatomical-part.Thumb* as an existential restriction in the definition of *Hand* would disallow to classify as *Hand* all those anatomical entities that have no thumb due to congenital or acquired abnormalities. As another example, most instances of *Ileum* do not contain a *Meckel's Diverticulum*, whereas all instances of *Meckel's Diverticulum* are necessarily *anatomical-part-of Ileum*. Many surgical interventions that remove anatomical structures (appendix, gallbladder, etc.), produce similar patterns.

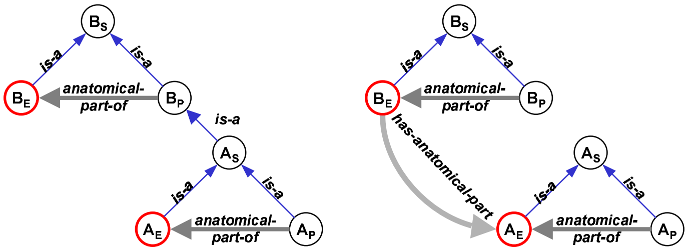

Figure 4: Patterns for Part-Whole Reasoning Using SEP triplets

In formalism we propose, this corresponds to a single taxonomic link between an S-node and a P-node (cf. Figure 4, left part). The contrary is also possible (cf. Figure 4, right part): the definition of A_E does not imply that the role *anatomical-part-of* be filled by B_E, but B_E does imply that the inverse role, *has-anatomical-part*, be filled by A_E. As an example, a *Lymph-node* necessarily contains *Lymph-follicles*, but there exist *Lymph-follicles* that are not part of a *Lymph-node*. This pattern is typical of the mereological relation between macroscopic (countable) objects, such as organs, and multiple uniform microscopic objects.

- *Analysis of the sibling relations.* In UMLS, the *SIB* relation links concepts that share the same parent in a taxonomic or partonomic hierarchy.

  Results: We found on the average 6.8 siblings per concept in the anatomy domain, 8.8 in the pathology domain. So far, the analysis of sibling relations has been performed only for the anatomy domain. From a total of 521 sibling relations, 9 were identified as *is-a*, 14 as *part-of*, and 17 as *has-part*, whereas 404 referred to topologically disconnected concepts.

- *Completion and modification of anatomy–pathology relations.* Surprisingly, only very few pathology concepts contained an explicit reference to a corresponding anatomy concept. Therefore, these relations have to be added by a domain expert. In each case, a decision must be made whether the E-node or the S-node has to be addressed as the target concept for modification such that the propagation of roles across part-whole hierarchies is disabled or enabled, respectively (cf. Figure 2).

  Results: In the sample, we found 522 anatomy-pathology relations, from which 358 (i.e., 69%!) were judged incorrect by the domain expert. In 36 cases an adequate anatomy-pathology relation was missing. All 164 *has-location* roles were analyzed as to whether they had to be filled by an S-node or an E-node of an anatomical triplet. In 153 cases, the S-node (which allows propagation

> across the part-whole hierarchy) was considered to be adequate, in 11 cases the E-node was preferred. The analysis of the 100 pathology concepts revealed that only 17 had to be linked with an anatomy concept. In 15 cases, the default linkage to the S-node was considered as correct, in one case the linkage to the E-node was preferred, in another case the linkage was considered as false.
>
> The high number of implausible constraints points to the lightweight semantics of *has-location* links in the UMLS sources. While we interpreted them as a conjunction for the mapping procedure, a disjunctive meaning seems to prevail implicitly in many definitions of top-level concepts such as *Tuberculosis*. In this example, *all* anatomical concepts that can be affected by this disease are linked by *has-location*. All these constraints (e.g., *has-location Urinary-Tract*) are inherited to subconcepts such as *Tuberculosis-of-Bronchus*. A thorough analysis of the top-level pathology concepts is necessary, and conjunctions of constraints will have to be substituted by disjunctions where necessary.

In conclusion, our study shows that it is relatively straightforward to restitute consistency of the UMLS knowledge base, but it is nearly impossible to reach a high degree of both adequacy and completeness due to the huge amount of manual work required. Restituting adequacy should, however, not be considered primarily as eliminating obvious categorization 'errors' from the UMLS sources, but rather as making choices between alternative conceptualizations of medical terms whose meaning differs slightly due to the heterogeneity of the knowledge sources. Another aspect is the need for the curation of concept definitions. These may have become obsolete as a consequence of imposing particularly rigid axiomatic assumptions on them. For instance, the conjunctive reading of defining attributes is not true in all cases, and, thus, necessarily requires individual manual specification. Finally, from a technical perspective, we found the implications of using the terminological classifier of utmost importance and of outstanding heuristic value. Hence, the knowledge refinement cycles are truly semi-automatic, fed by medical expertise on the side of the human knowledge engineer, but also driven by the reasoning system which makes explicit the consequences of (im)proper concept definitions.

References

1. A. McCray and S. Nelson. The representation of meaning in the UMLS. *Methods of Information in Medicine*, 34(1/2):193–201, 1995.
2. J. Cimino. Distributed cognition and knowledge-based controlled medical terminologies. *Artificial Intelligence in Medicine*, 12(1):153–168, 1998.
3. O. Bodenreider. Circular hierarchical relationships in the UMLS: Etiology, diagnosis, treatment, complications and prevention. In *Proc. AMIA 2001*, pages 57–61, 2001.

4. J. Cimino, P. Clayton, G. Hripsack, and S. Johnson. Knowledge-based approaches to the maintenance of a large controlled medical terminology. *JAMIA*, 1(1):35–50, 1994.
5. E. Mays, R. Weida, R. Dionne, M. Laker, B. White, C. Liang, and F. Oles. Scalable and expressive medical terminologies. In *Proceedings AMIA'96*, pages 259–263, 1996.
6. A. Rector, S. Bechhofer, C. Goble, I. Horrocks, W. Nowlan, and W. Solomon. The GRAIL concept modelling language for medical terminology. *Artificial Intelligence in Medicine*, 9:139–171, 1997.
7. F. Volot, M. Joubert, and M. Fieschi. Review of biomedical knowledge and data representation with Conceptual Graphs. *Methods of Inf. in Medicine*, 37(1):86–96, 1998.
8. A. Rector. Clinical terminology: Why is it so hard? *Meth. Inf. Med.*, 38:147–157, 1999.
9. D. Pisanelli, A. Gangemi, and G. Steve. An ontological analysis of the UMLS metathesaurus. In *Proceedings AMIA'98*, pages 810–814, 1998.
10. R. Côté. *SNOMED International*. College of American Pathologists, 1993.
11. K. Spackman and K. Campbell. Compositional concept representation using SNOMED: Towards further convergence of clinical terminologies. In *Proceedings AMIA'98*, pages 740–744, 1998.
12. K. A. Spackman. Normal forms for description logic expression of clinical concepts in SNOMED RT. In *Proc. AMIA 2001*, pages 627–631, 2001.
13. U. Hahn, M. Romacker, and S. Schulz. Creating knowledge repositories from biomedical reports: The MEDSYNDIKATE text mining system. In *Proceedings PSB 2002*, pages 338–349, 2002.
14. A. Artale, E. Franconi, N. Guarino, and L. Pazzi. Part-whole relations in object-centered systems: An overview. *Data & Knowledge Engineering*, 20(3):347–383, 1996.
15. I. Horrocks and U. Sattler. A description logic with transitive and inverse roles and role hierarchies. *Journal of Logic and Computation*, 9(3):385–410, 1999.
16. U. Hahn, S. Schulz, and M. Romacker. Part-whole reasoning: A case study in medical ontology engineering. *IEEE Intelligent Systems & Applications*, 14(5):59–67, 1999.
17. D. Cruse. On the transitivity of the part-whole relation. *J. Ling.*, 15:29–38, 1979.
18. M. Winston, R. Chaffin, and D. Herrmann. A taxonomy of part-whole relationships. *Cognitive Science*, 11:417–444, 1987.
19. A. Rector. Analysis of propagation along transitive roles: Formalisation of the GALEN experience with medical ontologies. In *Proceedings DL 2002*, 2002.
20. J. Schmolze and W. Mark. The NIKL experience. *Comput. Intell.*, 7(1):48–69, 1991.
21. E. Schulz, C Price, and P. Brown. Symbolic anatomic knowledge representation in the Read Codes Version 3: Structure and application. *JAMIA*, 4(1):38–48, 1997.
22. M. Schmidt-Schauß and G. Smolka. Attributive concept descriptions with complements. *Artificial Intelligence*, 48(1):1–26, 1991.
23. S. Schulz and U. Hahn. Necessary parts and wholes in bio-ontologies. In *Proceedings KR 2002*, pages 387–394, 2002.
24. R. MacGregor and R. Bates. The LOOM knowledge representation language. Technical Report RS-87-188, Information Sciences Institute, Univ. of Southern California, 1987.
25. J. Heinsohn, D. Kudenko, B. Nebel, and H.-J. Profitlich. An empirical analysis of terminological representation systems. *Artificial Intelligence*, 68(2):367–397, 1994.