


Towards Very Large Ontologies for Medical Language Processing

Udo Hahn Stefan Schulz

 Text Knowledge Engineering Lab

Freiburg University

Werthmannplatz 1

D-79098 Freiburg, Germany

<http://www.coling.uni-freiburg.de/>

Abstract

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. Our approach consists of four steps: concept definitions are automatically generated from the UMLS source, integrity checking of taxonomic and paronomic hierarchies is performed by the terminological classifier, cycles and inconsistencies are eliminated, and incremental refinement of the evolving knowledge base is performed by a domain expert. We report on experiments with a knowledge base composed of 164,000 concepts and 76,000 relations.

1. Introduction

Medical language processing (for a survey, cf. Spyns (1996)) is an application area of natural language processing, which deals with the analysis of medical documents such as discharge, X-ray, admission, and finding reports. The sublanguage patterns these documents exhibit reveal that understanding their contents, whether by humans or machines, requires access to knowledge and the terminology of the medical domains being covered.

Unlike many other disciplines, medicine has a long standing tradition in assembling and structuring medical knowledge, e.g. disease taxonomies, medical procedures, anatomical terms, etc., 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 (quasi-)synonymous terms. This is most evident in the UMLS, the *Unified Medical Language System* (McCray and Nelson, 1995; McCray, 1998), an umbrella system which covers more than 60 medical thesauri and classifications. Its metathesaurus component contains more than 800,000 concepts which are structured in hierarchies and classified by 134 semantic types (provided by the UMLS *Semantic Network*). Their semantics is shallow and entirely intuitive, which is due to the fact that their usage was primarily intended for humans as a backbone for various forms of clinical knowledge management, e.g., cross-mapping between different terminologies, disease and procedure encoding, etc.

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. (Cimino, 1998). This may not cause utterly severe problems when humans are in the loop and its use is limited to disease encoding, accountancy or document retrieval tasks. However, anticipating its use for more knowledge-intensive applications such as natural language understanding of medical narratives (Hahn et al., 1999b) those shortcomings might lead to an impasse.

As a consequence, formal models for dealing with medical knowledge have been proposed, using representation mechanisms based on conceptual graphs, semantic net-

works or description logics (Cimino et al., 1994; Mays et al., 1996; Rector et al., 1997; Volot et al., 1998; Gangemi et al., 1999). Not surprisingly, there is a price to be paid for more expressiveness and formal rigor, *viz.* increasing modeling efforts and, hence, increasing maintenance costs. Operational systems making full use of this rigid approach, 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 (Rector et al., 1997). The limited coverage then hampers their routine usage, an issue which is always highly rewarded in the medical informatics community.

The knowledge bases developed within the framework of the above-mentioned terminological systems have almost all been designed from scratch – without making systematic use of the large body of knowledge contained in those medical terminologies. An intriguing approach would be to join the massive *coverage* offered by informal medical terminologies with the high level of *expressiveness* supported by rigid knowledge representation systems in order to develop formally solid medical knowledge bases on a larger scale. In the paper, we describe such a knowledge engineering methodology. The resulting medical ontology forms the domain knowledge backbone of MEDSYNDIKATE, a system for the automatic acquisition of factual and evaluative (subjective) knowledge from medical finding reports (Hahn et al., 1999b; Hahn et al., 1999a). Its size, more than 240,000 concepts and roles, is truly competitive and supplies the system with a fair amount of medical knowledge in the core areas of anatomy and pathology.

2. Reasoning Along Part-Whole Hierarchies

Medical ontologies are organized around taxonomic (IS-A relation) and paronomic (PART-OF relation) knowledge. Hence, medical knowledge representation efforts have to take account of both hierarchy types and the reasoning patterns they imply.¹

¹In an exploratory quantitative study of the particular text genre we are dealing with, *viz.* medical finding reports, we found

Partonomic knowledge has been an issue within diverse areas ranging from philosophy (mereology), data modeling for database systems and object-oriented programming, to knowledge representation proper in the field of artificial intelligence. Major strands of this work are discussed by Artalet et al. (1996) under the heading of object-centered representation approaches. This also includes the description logic (DL) paradigm to which we subscribe in our work, too. From our application domain, the need arises to have formally solid inference mechanisms for taxonomic (generalization hierarchies), as well as partonomic reasoning (part-whole hierarchies) available within a uniform representation model. We also require an inference engine which performs this style of advanced reasoning on large data sets ($\gg 10,000$ items). Hence, we consider KL-ONE-type descriptions logics (Woods and Schmolze, 1992), at the formal representation level, and LOOM's classification-based inference machine (MacGregor and Bates, 1987; MacGregor, 1994), at the system level, the most convenient match of our requirements and the current state of the art in terminological reasoning.

Unlike generalization-based reasoning in concept taxonomies, no fully conclusive mechanism exists up to now for reasoning along part-whole hierarchies. In the description logics community several language extensions have been proposed which provide special constructors for part-whole reasoning (Rector et al., 1997; Horrocks and Sattler, 1999). This seems a reasonable way to proceed, as long as the transitivity property of a relation can be assumed, in general. In the medical (Hahn et al., 1999c) as well as commonsense domains (Cruse, 1979), however, 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. Motivated by informal approaches sketched by Schmolze and Mark (1991) and Schulz et al. (1997), we formalized a model of partonomic reasoning (Hahn et al., 1999c) that meets the above requirements and also does not exceed the expressiveness of the well-understood, parsimonious concept language \mathcal{ALC} (Woods and Schmolze, 1992).²

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.

(Hahn et al., 1999a) that 45% of the (indirect) reference relations between sentences are due to bridging or functional anaphora, which refer heavily to partonomic knowledge. On the other hand, 34% of the (direct) reference relations are due to nominal anaphora, which bring taxonomic knowledge to bear.

² \mathcal{ALC} allows for the construction of hierarchies of concepts and relations, where \sqsubseteq denotes subsumption and \doteq definitional equivalence. Existential (\exists) and universal (\forall) quantification, negation (\neg), disjunction (\sqcup) and conjunction (\sqcap) are supported. Role filler constraints (e.g., typing by C) are linked to the relation name R by a dot, $\exists R.C$.




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). Each *Structure* concept subsumes directly both an anatomical *Entity* and each of the anatomical *Parts* of this entity, i.e., the corresponding **E-node** and **P-node**, e.g., HAND and HAND-PART, respectively. Unlike entities and their parts which refer to specific ontological objects, structures have no physical correlate in the real world — they constitute a representational artifact required for the formal reconstruction of systematic patterns of partonomic reasoning. The E-node denotes the whole anatomical entity to be modeled, whereas the 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. Hence, for every P-node there exists a corresponding E-node for the role ANATOMICAL-PART-OF. A reconstruction of some basic anatomical relations in terms of SEP triplets is illustrated in Figure 2.




Figure 2: SEP Triplet Model of a Partonomic Hierarchy

The reconstruction of the relation ANATOMICAL-PART-OF by taxonomic reasoning proceeds as follows. Let us assume that C_E and D_E denote E-nodes, C_S and D_S denote the S-nodes that subsume C_E and D_E , respectively, and C_P and D_P 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:

$$C_E \sqsubseteq C_S \sqsubseteq D_P \sqsubseteq D_S \quad (1)$$

$$D_E \sqsubseteq D_S \quad (2)$$

The P-node is defined as follows (note the disjointness between D_E and D_P , i.e., no anatomical concept can be ANATOMICAL-PART-OF itself):

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

		Step 1	Step 2	Step 3	Step 4				
UMLS relation	number of links	Automatic generation of Loom definitions, augmented by P-Loom language elements ;;; = comment line	Submission to Loom classifier. Validation for formal consistency by Loom	Manual restitution of formal consistency	Manual rectification and refinement of the resulting knowledge base				
Anatomy Concepts Linked to Anatomy Concepts									
sibling_of	267.218	;;; SIB			add negations in order to express taxonomic or partitive disjointness				
child_of	59.808	;;; CHDRN			include related concepts into :is-primitive or :part-of clause where plausible				
narrower_term	24.223	;;; CHDRN							
isa	9.755	:is-primitive	check for definitional cycles	remove taxonomic parent concepts	substitute of primitive links by non-primitive ones where possible				
location_of	4.803	;;; LOCATION_OF			include related concepts into :has-part clause where plausible				
has_location	4.803	;;; HAS_LOCATION			include related concepts into :part-of clause, where plausible				
has_part	4.321	:has-part			check whether this part is mandatory (under "real-anatomy" assumption)				
has_conceptual_part	126								
part_of	4.321	:part-of	1. check for partonomic cycles 2. check for disjointness between E and P node	1. remove partonomic or taxonomic parent concepts 2. redefine triplet as single concept	check for plausibility and completeness				
conceptual_part_of	126								
parent	59.808	;;; PARRB			include related concepts into :has-part clause where plausible				
broader_term	24.223	;;; PARRB							
inverse_isa	9.755	<do nothing>							
associated_with	14								
mapped_from	2.643								
other_relation	10.908								
qualified_by	1.864								
allowed_qualifier	1.864								
mapped_to	2643								
<other named relations>	11.886					(:some x)	check for inherited constraints	remove constraints	remove or add constraints
Pathology Concepts Linked to Pathology Concepts									
sibling_of	457.542					;;; SIB			add negations in order to express taxonomic disjointness
child_of	72.426	:is-primitive	check for definitional cycles	remove parent concepts	substitute primitive links by non-primitive ones whenever possible				
narrower_term	26.972								
isa	3.635								
inverse_isa	3.635								
associated_with	13.902								
mapped_to	15.024								
mapped_from	15.024								
part_of	1								
has_part	1					<do nothing>			
parent	72.426								
broader_term	28.972	(:some x)	check for inherited constraints	remove constraints	remove or add constraints				
other_relation	25.796								
qualified_by	6.255								
allowed_qualifier	6.255								
<other named relations>	4.162								
Pathology Concepts Linked to Anatomy Concepts									
CUIpat = CUIana	2.247	(:some has_anatomic_correlate)			plausibility check of concept "duplication" (assignment to both domains)				
<missing>		<do nothing>			add pathology-anatomy links				
associated_with	2.314	(:some associated_with <anatomy_concept> S)		check for consistency	render links complete, link to E-node instead of S-node when role propagation has to be disabled				
has_location	9,230	(:some has_location <anatomy_concept> S)							
<other>		<do nothing>							

Figure 4: Workflow Diagram for the Construction of a LOOM Knowledge Base from the UMLS

(deftriplet HEART :is-primitive HOLLOW-VISCUS :has-part (:p-and FIBROUS-SKELETON-OF-HEART WALL-OF-HEART CAVITY-OF-HEART LEFT-SIDE-OF-HEART RIGHT-SIDE-OF-HEART AORTIC-VALVE PULMONARY-VALVE))
--

Table 1: Generated Triplets in LOOM Format

For the pathology domain, we treated *CHD* (child) and *RN* (narrower relation) from the UMLS as indicating taxonomic links. No part-whole relations were considered, since this category does not apply to the pathology domain. 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 semantic assumption all roles generated in this process were considered as 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 role $R.B$ 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.

In both subdomains, shallow relations such as the extremely frequent sibling relation (*SIB*) were included as comments into the code to provide heuristic guidance for the subsequent manual refinement phase.

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 semantic network) a total of 240,764 definition 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 (average cardinality: 3.3). 4,043 DEFTRIPLET statements contained a :PART-OF clause, only in 332 cases followed by more than one argument (average cardinality: 1.1). The resulting knowledge base was then submitted to the terminological classifier 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 (cf. Table 2).

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*

	Anatomy	Pathology
Triplets	38,059	—
defconcept statements	114,177	50,087
cycles	1	355
inconsistencies	2,328	0

Table 2: Classification Results

links, an encoding that raises a conflict due to the disjointness required for corresponding P- and E-nodes. In most of these cases the affected parents 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 involved triplets with simple LOOM concepts, by matching them with already existing P-nodes or by disabling IS-A or PART-OF links.

In the pathology part of the knowledge base, we expected a large number of terminological cycles, as a consequence of interpreting the thesaurus-style *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 a tolerable number. 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 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.

The cycles were analyzed and a negative list which consisted of 630 concept pairs was manually derived. In a subsequent extraction cycle we incorporated this list in the automated construction of the LOOM concept definitions and, with these new constraints, a fully consistent knowledge base was generated.

Step 4: Manual Rectification and Refinement 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, in total. The fourth step – when performed for the whole knowledge base – is very time-consuming and requires broad and in-depth medical expertise. An analysis of random samples from both subdomains is currently being performed by the second author, a domain expert. The preliminary 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 in the anatomy and pathology subdomains 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. Primitive subsumption is substituted by non-primitive one whenever possible. This is a crucial point, because the automatically generated hierarchies contain only information about the parent concepts and necessary conditions.

As an example, the automatically generated definition of DERMATITIS includes the information that it is an INFLAMMATION, and that the role HAS-LOCATION must be filled by the concept SKIN. An INFLAMMATION that HAS-LOCATION SKIN, however, cannot automatically be classified as DERMATITIS.

Results: In the *anatomy* sample, only 76 concepts out of 100 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 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 being subsumed by 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 concepts that may exhibit pathological modifications. Figure 6 (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.⁴ As an example,




Figure 6: Patterns for Part-whole Reasoning Using SEP triplets

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. In our formalism, this corresponds to a single taxonomic link between an S-node and a P-node (cf. Figure 6, left part). The contrary is also possible (cf. Figure 6, 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 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 and defining concepts as being disjoint.* In UMLS, the *SIB* relation links concepts that share the same parent in a taxonomic or partonomic hierarchy. Pairs of sibling concepts may have common descendants or not. If not, they constitute the root of two disjoint subtrees. In a taxonomic hierarchy, this means that one concept implies the negation of the other (e.g., a benign tumor cannot be a malignant one, *et vice versa*). In a partitive hierarchy, this can be interpreted as spatial disjointness, *viz.* one concept does not spatially overlap with another one. As an example, ESOPHAGUS and DUODENUM are spatially disjoint, whereas STOMACH and DUODENUM are not (they share a common transition structure, called PYLORUS), such as all neighbor structures that have a surface or region in common. Spatial disjointness can be modeled so that the definition of the S-node of the concept A implies the negation of the S-node of the concept B (Schulz et al., 2000).

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.

VALVE and PULMONARY-VALVE should be eliminated from the :HAS-PART list, because they may be missing in certain cases as a result of congenital malformations, inflammatory processes or surgical interventions.

⁴In Table 1 the concepts marked by *italics*, *viz.* AORTIC-

- *Completion and modification of anatomy–pathology relations.* Surprisingly, only very few pathology concepts contained an explicit reference to a corresponding anatomy concept. These relations must, therefore, be added by a domain expert. In each case, the 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.

Results: In the sample we found 522 anatomy-pathology relations, from which 358 (69%!) were judged incorrect by the domain experts. In 36 cases an adequate anatomy-pathology relation was missing. All 164 HAS-LOCATION roles were analyzed as to whether they were 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 were to be linked with an anatomy concept. In 15 cases, the default linkage to the S-node was considered to be correct, in one case the linkage to the E-node was preferred, in another case the linkage was considered to be false.

The high number of implausible constraints points to the lightweight semantics of HAS-LOCATION links in the UMLS sources. While we interpreted them in terms of a conjunction for the import routine, a disjunctive meaning seems to prevail implicitly in many definitions of top-level concepts such as TUBERCULOSIS. In this example, we find all anatomical concepts that can be affected by this disease linked by HAS-LOCATION. All these constraints (e.g., HAS-LOCATION URINARY-TRACT) are inherited to sub-concepts 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.

4. Discussion and Conclusions

In medical language processing, knowledge of the underlying domain has to be supplied on a larger scale. Instead of developing sophisticated medical knowledge bases from scratch, we here propose a ‘conservative’ approach — reuse existing large-scale resources, but refine the data from these resources so that advanced representational requirements imposed by more expressive knowledge representation languages are met. The resulting knowledge bases can then be used for sophisticated applications requiring formally sound medical reasoning such as text understanding.

The benefits and problems of converting conceptual knowledge from semantically weak specifications to a rigorous knowledge representation formalism have been described by Pisanelli et al. (1998). They extracted knowledge from the UMLS semantic network, as well as from parts of the metathesaurus and converted it into a description logics system. Spackman and Campbell (1998) describe how the SNOMED nomenclature evolves from a multi-axial coding system into a formally founded ontol-

ogy. Their general goal is to avoid ambiguous or semantically invalid representations of composite concepts. However, both approaches do not provide a special reasoning mechanism for partonomic relations.

Within the formal framework of GALEN, a fragment of the Read Thesaurus was translated into GRAIL, a knowledge representation system also based on description logics (Rogers et al., 1998). In a cross-validation study it was checked, on the one hand, whether the definitions contained in the Read Thesaurus were logically consistent and, on the other hand, whether the GRAIL domain model was rich enough to encode them. Although GRAIL comes with a special-purpose reasoning mechanism dedicated to partonomies, the adaptation was limited to simple generic hierarchies as only these structure the Read Thesaurus.

The developers of VOXEL-MAN (Schubert and Höhne, 1998), a multimedia tutoring systems for anatomy, and of the Digital Anatomist (UWDA), an anatomical semantic network (Rosse et al., 1998), have both emphasized partitive hierarchies though at an informal level. Whereas in VOXEL-MAN a fine-grained ontology of partonomic relations is sketched that accounts for various part-whole relations found in the anatomy domain, the UWDA developers restrict themselves to a small set of relations leading to a precise separation between partonomic and taxonomic hierarchies. They excel with a high granularity of description and a broad coverage.

Our approach tries to combine the broad coverage and fine-grained concept descriptions of the UWDA with the formal rigor of description logics. Additionally, we enhance the imported knowledge with part-whole specific reasoning capabilities indispensable in the medical domain, though this has already been described as a hard problem for terminological languages (Haimowitz et al., 1988).

Several alternatives have been brought forward to solve this problem within the DL framework. The first approach extends a basic language definition (usually, *ALC*) by generally allowing role hierarchies and transitivity of roles (Horrocks and Sattler, 1999), or by supplying dedicated transitive partonomy operators such as for the GRAIL language (Rector et al., 1997). This way, the transitivity property is hard-wired to role definitions and assumed to hold, in general. This contradicts empirical data from anatomical ontologies as discussed in Hahn et al. (1999c), which reveal that the same relation may be transitive and nontransitive depending on the anatomical context.

Alternatively, one may want to preserve standard language definitions for reasons of simplicity and parsimony. Accordingly, Schmolze and Mark (1991) propose a ‘data structure’ solution similar to ours using subsumption to obtain inferences resembling those of transitive roles or transitive closure of roles. Artale et al. (1996) criticize this proposal for the “proliferation of (artificial) concepts” involved. We argue, however, that these additional concepts are necessary from an ontological point of view, as the distinct mechanisms for conditioned specialization modeling reveal (cf. Figure 3).

It remains to be seen whether conservative structural extensions of a stable language platform are able to carry over to the many varieties of partonomic reasoning and differ-

ent part-whole relations, or whether newly designed operators or other fundamental language extensions are needed. In the medical domain, at least, where the restriction to one subrelation of PART-OF, viz. ANATOMICAL-PART-OF, is sufficient, a relatively simple “data structure” extension like the SEP triplets yields already adequate results, without the necessity to resort to profound extensions of the terminological language. We have evidence that the triplet mechanism we here propose can be straightforwardly extended to cover mereotopological and (limited) spatial reasoning, as well (Schulz et al., 2000; Schulz and Hahn, 2001).

5. References

- A. Artale, E. Franconi, N. Guarino, and L. Pazzi. 1996. Part-whole relations in object-centered systems: An overview. *Data & Knowledge Engineering*, 20:347–383.
- J. Cimino, P. Clayton, G. Hripsack, and S. Johnson. 1994. Knowledge-based approaches to the maintenance of a large controlled medical terminology. *Journal of the American Medical Informatics Association*, 1(1):35–50.
- J. Cimino. 1998. Distributed cognition and knowledge-based controlled medical terminologies. *Artificial Intelligence in Medicine*, 12(1):153–168.
- D. Cruse. 1979. On the transitivity of the part-whole relation. *Journal of Linguistics*, 15:29–38.
- A. Gangemi, D. Pisanelli, and G. Steve. 1999. An overview of the ONION project: Applying ontologies to the integration of medical terminologies. *Data & Knowledge Engineering*, 31(2):183–220.
- U. Hahn, M. Romacker, and S. Schulz. 1999a. Discourse structures in medical reports – watch out! The generation of referentially coherent and valid text knowledge bases in the MEDSYNDIKATE system. *International Journal of Medical Informatics*, 53(1):1–28.
- U. Hahn, M. Romacker, and S. Schulz. 1999b. How knowledge drives understanding: Matching medical ontologies with the needs of medical language processing. *Artificial Intelligence in Medicine*, 15(1):25–51.
- U. Hahn, S. Schulz, and M. Romacker. 1999c. Part-whole reasoning: A case study in medical ontology engineering. *IEEE Intelligent Systems & Their Applications*, 14(5):59–67.
- I. Haimowitz, R. Patil, and P. Szolovits. 1988. Representing medical knowledge in a terminological language is difficult. In *SCAMC’88 – Proceedings of the 12th Annual Symposium on Computer Applications in Medical Care*, pages 101–105.
- I. Horrocks and U. Sattler. 1999. A description logic with transitive and inverse roles and role hierarchies. *Journal of Logic and Computation*, 9(3):385–410.
- R. MacGregor and R. Bates. 1987. The LOOM knowledge representation language. Technical Report RS-87-188, Information Sciences Institute, USC.
- R. MacGregor. 1994. A description classifier for the predicate calculus. In *AAAI’94 – Proceedings 12th National Conference on Artificial Intelligence*, pages 213–220.
- E. Mays, R. Weida, R. Dionne, M. Laker, B. White, C. Liang, and F. Oles. 1996. Scalable and expressive medical terminologies. In *AMIA’96 – Proceedings of the AMIA Annual Fall Symposium*, pages 259–263.
- A. McCray and S. Nelson. 1995. The representation of meaning in the UMLS. *Methods of Information in Medicine*, 34(1/2):193–201.
- A. McCray. 1998. The nature of lexical knowledge. *Methods of Information in Medicine*, 37(4/5):353–360.
- D. Pisanelli, A. Gangemi, and G. Steve. 1998. An ontological analysis of the UMLS metathesaurus. In *AMIA’98 – Proceedings of the AMIA Annual Fall Symposium*, pages 810–814.
- A. Rector, S. Bechhofer, C. Goble, I. Horrocks, W. Nowlan, and W. Solomon. 1997. The GRAIL concept modelling language for medical terminology. *Artificial Intelligence in Medicine*, 9:139–171.
- J. Rogers, C. Price, A. Rector, W. Solomon, and N. Smeijko. 1998. Validating clinical terminology structures: Integration and cross-validation of READ THESAURUS and GALEN. In *AMIA’98 – Proceedings of the AMIA Annual Fall Symposium*, pages 845–849.
- C. Rosse, J. Mejino, B. Modayur, R. Jakobovits, K. Hinshaw, and J. Brinkley. 1998. Motivation and organizational principles for anatomical knowledge representation: The DIGITAL ANATOMIST symbolic knowledge base. *Journal of the American Medical Informatics Association*, 5(1):17–40.
- J. Schmolze and W. Mark. 1991. The NIKL experience. *Computational Intelligence*, 6(1):48–69.
- R. Schubert and K.-H. Höhne. 1998. Paronomies for interactive explorable 3D-models of anatomy. In *AMIA’98 – Proceedings of the AMIA Annual Fall Symposium*, pages 433–437.
- S. Schulz and U. Hahn. 2001. Mereotopological reasoning about parts and (w)holes in bio-ontologies. In C. Welty and B. Smith, editors, *Formal Ontology in Information Systems. Collected Papers from the 2nd International Conference*, pages 210–221.
- E. Schulz, C. Price, and P. Brown. 1997. Symbolic anatomic knowledge representation in the Read Codes Version 3: Structure and application. *Journal of the American Medical Informatics Association*, 4(1):38–48.
- S. Schulz, U. Hahn, and M. Romacker. 2000. Modeling anatomical spatial relations with description logics. In *AMIA 2000 – Proceedings of the AMIA Annual Fall Symposium*, pages 779–783.
- K. Spackman and K. Campbell. 1998. Compositional concept representation using SNOMED: Towards further convergence of clinical terminologies. In *AMIA’98 – Proceedings of the AMIA Annual Fall Symposium*, pages 740–744.
- P. Spyns. 1996. Natural language processing in medicine: An overview. *Methods of Information in Medicine*, 35(4/5):285–301.
- F. Volot, M. Joubert, and M. Fieschi. 1998. Review of biomedical knowledge and data representation with Conceptual Graphs. *Methods of Information in Medicine*, 37(1):86–96.
- W. Woods and J. Schmolze. 1992. The KL-ONE family. *Computers & Mathematics with Applications*, 23(2/5):133–177.