Context-Based Task Ontologies for Clinical Guidelines

Anand KUMAR[a,b], Paolo CICCARESE[a], Barry SMITH[b,c], Matteo PIAZZA[a]

[a] *LMI, Department of Computer Science, University of Pavia, Italy*
[b] *IFOMIS, University of Leipzig, Germany*
[c] *Department of Philosophy, University at Buffalo, New York, USA*

Abstract. Evidence-based medicine relies on the execution of clinical practice guidelines and protocols. A great deal of effort has been invested in the development of tools which can automate the representation and execution of the recommendations contained within such guidelines, by creating Computer Interpretable Guideline Models (CIGMs). Context-based task ontologies (CTOs), based on standard terminology systems like UMLS, form one of the core components of such models. We have created DAML+OIL-based CTOs for the tasks referred to in the WHO guideline for hypertension management, drawing comparisons also with other, related guidelines. The advantages of CTOs include: contextualization of ontologies, tailoring of ontologies to specific aspects of the phenomena of interest, division of the complex tasks involved in creating ontologies into different levels, and provision of a methodology by means of which the task recommendations contained within guidelines can be integrated into the clinical practices of a health care set-up.

1. Introduction

1.1 Clinical Practice Guidelines

Clinical Practice Guidelines (CPGs) are 'systematically developed statements to assist practitioner and patient decisions about appropriate health care for specific clinical circumstances' [1]. Their use in clinical decision-making is intended to improve the outcomes of clinical care [2]. Given that most CPGs are a formulated as unstructured text or as simple flowcharts, there is a growing need to create Computer Interpretable Guideline Models (CIGMs). For this, however, we require standardized terminologies based on coherent ontologies of clinical activities.

1.2. The UMLS Semantic Network and Metathesaurus

The Unified Medical Language System (UMLS), designed by the National Library of Medicine, integrates a number of standard medical terminologies into a single unified knowledge representation system [3]. The UMLS Semantic Network consists of 134 Semantic Types and 54 links between these types. These form a graph with a double tree structure, the topmost nodes being '*Entity*' and '*Event*'. The vertices of this graph are the Semantic Types, the edges are the links between them. The corresponding complete graph would contain more than 6000 edges. Even this, however, represents merely a convenient high-level abstraction from the entire UMLS, whose concept repository, the UMLS Metathesaurus (META), includes as of January 2003 some 875,255 concepts and 2.14 million concept names.

1.3 DAML + OIL

While the UMLS provides the terms contained within it with associated Semantic Types, one needs a more elaborate ontology in order to use the latter in CIGMs. Among the emerging standards in this field, the DARPA Agent Markup Language and Ontology Interface Language (DAML+OIL) [4] is a proposal for an ontology representation language suitable for these purposes. DAML+OIL is complemented by OilEd, an ontology editor that supports the construction of OIL-based ontologies.

2. The UMLS Semantic Network, DAML+OIL and Clinical Practice Guidelines

2.1 UMLS Semantic Types for Task-Based Clinical Practice Guidelines

UMLS ST	Definition
Diagnostic Procedure	A method, procedure or technique used to determine the nature or identity of a disease or disorder. This excludes procedures which are primarily carried out on specimens in a laboratory
Laboratory Procedure	A procedure method or technique used to determine the composition, quality, or concentration of a specimen, and which is carried out in a clinical laboratory. Included here are procedures which measure the times and rates of reactions
Therapeutic or Preventive Procedure	A procedure method or technique designed to prevent a disease or a disorder, or to improve physical function, or used in the process of treating a disease or injury

Table 1 Definitions of the 3 UMLS Semantic Types used in the task-based CIGMs.

[5] shows on the basis of manual mark-up of CPG texts that most of the actions suggested in CPGs can be mapped into that part of the UMLS terminology which is associated with one or other of the Semantic Types *Laboratory Procedure*, *Diagnostic Procedure* and *Therapeutic or Preventive Procedure* (see Table 1), all of which are subclassifications of the Semantic Type *Health Care Activity*.

Other Semantic Types closely associated with *Health Care Activity* but used less frequently in CPGs are: *Educational Activity*, *Governmental or Regulatory Activity* and *Research Activity*, all of which are subclassifications of *Occupational Activity*. *Research Activity*, for example, helps to determine an instance of *Health Care Activity* by contributing its "strength of evidence".

2.2 DAML+OIL-Based Ontology of a simplified UMLS Semantic Network

OilEd [6, 7], the DAML+OIL editor, was used in order to specify the relationships between our three selected Semantic Types and the remaining Semantic Types of UMLS. Among all possible relationships (edges), we selected all those adjacent to our three initial vertices together with all edges immediately adjacent to these. In this way we were able to form from

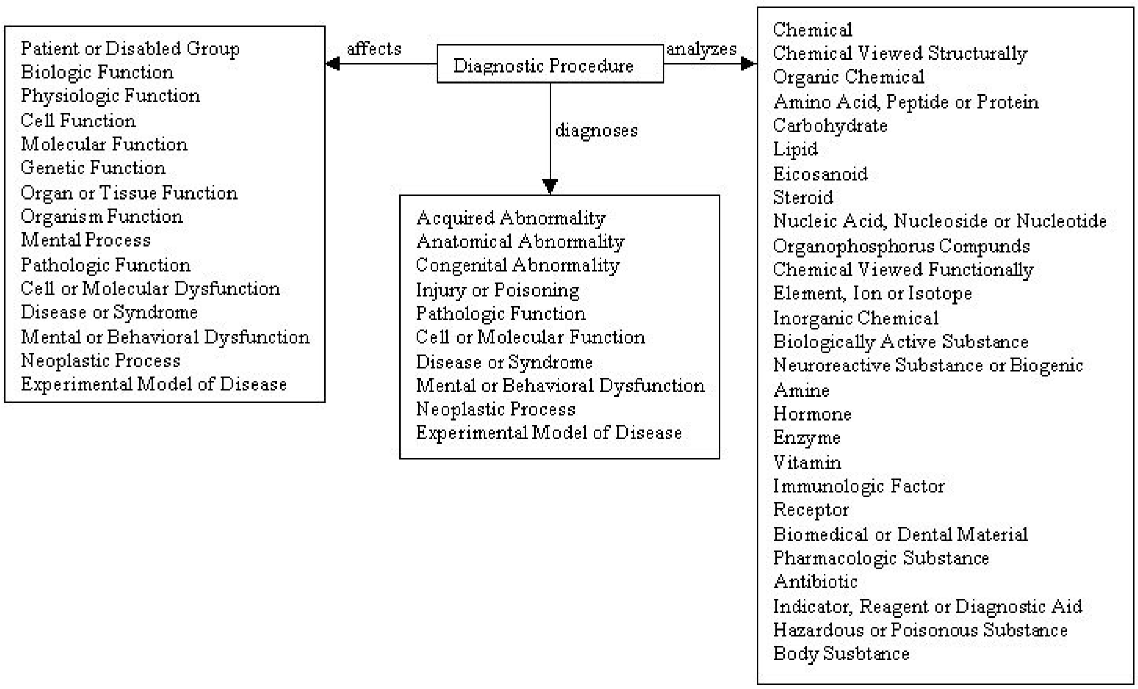

Figure 1. Some adjacency relations of Diagnostic Procedure in UMLS.

the original graph of the UMLS Semantic Network a Minimal Spanning Subgraph which includes all the original vertices but decreases the number of edges to the minimum needed to make a connected subgraph. The adjacency relationships for *Diagnostic Procedure* are shown in Figure 1.

In cases where the relations involved are semantically similar – for example "analyzes" and "diagnoses" – they were combined together to reduce complexity. Since we are focusing in our ontologies exclusively on *tasks*, we were able to effect a further simplification by using modifiers such as 'determination of' (DOF) as a means for converting classes in other Semantic Types into *Diagnostic Procedures*. For example "Proteinuria" (Figure 2), is either a laboratory or test result or a disease or syndrome. None of these is a *Health Care Activity*. Hence we use the DOF operator in order to incorporate "Proteinuria" into our CPG ontology (which is restricted to the Semantic Type *Health Care Activity*) via the heading: "Determination of Proteinuria" (DOF Proteinuria).

Term – Proteinuria

Semantic Type – *Laboratory or Test Result, Disease or Syndrome*

Operator – DOF (Determination of)

Term – DOF Proteinuria

New Semantic Type – *Laboratory Procedure*

Figure 2. Mapping for the term “Proteinuria”

Term – Proteinuria

Semantic Type – *Laboratory or Test Result, Disease or Syndrome*

Operator – DOF (Determination of)

Term – DOF Proteinuria

New Semantic Type – *Laboratory Procedure*

Figure 3. Mapping for the term “Proteinuria”

2.3 DAML+OIL Ontology of Laboratory Procedures

Our design methodology was, first, to construct a basic framework of classes and slots for the CPG domain, starting with the two UMLS superclasses ‘*Entity*’ and ‘*Event*’, and then to incrementally extend and refine this framework by adding new classes and slots together with

. yzes

tory

Procedure is a subclass of the class of *Health Care Activities* and a *Chemical* is a subclass of the class of *Substances*; a *Laboratory Procedure* analyzes a *Chemical* and a *Chemical* is analyzed by a *Laboratory Procedure*:

class-def defined *Laboratory Procedure*

subclass-of *Health Care Activity*

slot-constraint analyzes

value-type *Chemical*

class def defined *Chemical*

subclass-of *Substance*

slot-def analyzes

inverse is analyzed by

Figure 3. Frame based approach for Laboratory Procedure that analyzes Chemicals

However, the problem with the frame-based approach is that it can specify only certain relations between the types of classes involved. Thus, representing *Laboratory Procedure* as a subclass of *Health Care Activity* does not give us any insight into whether *Laboratory Procedure* is the only subclass of *Health Care Activity* and whether completion of the former would also mean completion of the latter.

These issues can be addressed by adding to our ontologies some of the resources of mereology or the theory of relations between parts and wholes ([8]). We depict *Laboratory Procedure* as a part-of *Health Care Activities involving Laboratory Procedure* (that is each instance of *Laboratory Procedure* is a part of some instance of *Health Care Activities involving Laboratory Procedure)*. This approach allows us to specify whether *Laboratory Procedure* is or is not the only part-of a given *Health Care Activity*. In other words it allows us to represent the fact that *Health Care Activity – Laboratory Procedure* is or is not empty, where '–' symbolizes the mereological subtraction operator. (For the formal treatment of these matters see [9], [10].)

Adding such mereological resources amounts to creating a more fine-grained partition of the reality represented in any ontology. Partitions are used to create refined classifications, for example, *Health Care Activity* can be classified into two categories by a partition which divides health care activities into those involving and those not involving some *Laboratory Procedure*. A mereologized partition might be represented in DAML+OIL as follows:

```
class-def defined Laboratory Procedure
        subclass-of Health Care Activity
        slot-constraint part-of
                value-type Health Care Activities involving Laboratory Procedures
        slot-constraint analyzes
                value-type Chemical

class-def defined Diagnostic Procedure
        subclass-of Health Care Activity
        slot-constraint part-of
                value_type Health Care Activities involving Diagnostic Procedure

class-def defined Therapeutic and Preventive Procedure
        subclass-of Health Care Activity
        slot-constraint part-of
                value-type Health Care Activities involving Therapeutic and Preventive Procedure

class-def defined Health Care Activity
        slot-constraint part-of
                value-type Health Care Activities involving Laboratory Procedures
        slot-constraint part-of
                value-type Health Care Activities involving Diagnostic Procedure
        slot-constraint part-of
                value-type Health Care Activities involving Therapeutic and Preventive Procedure

slot-def analyse
        inverse is-analyzed_by


slot-def part-of
        inverse has_part

disjoint
        Laboratory Procedure
        Diagnostic Procedure
        Therapeutic and Preventive Procedure

class-def defined Health Care Activities

        disjointUnionOf

                Health Care Activities involving Laboratory Procedures

                Health Care Activities involving Diagnostic Procedure

                Health Care Activities involving Therapeutic and Preventive Procedure
```

Figure 4. Mereological Partition of Semantic Types relevant to Clinical Guideline ontologies

It is important to notice that the last construct determines an exhaustive partition in the sense that it ensures that *Health Care Activities* is composed by and only by: *Health Care Activities involving Laboratory Procedure*, *Health Care Activities involving Diagnostic Procedure* and *Health Care Activities involving Therapeutic and Preventive Procedure*. Moreover, it ensures that *Health Care Activities involving Laboratory Procedures* includes at least a *Laboratory Procedure* but can contain also other kinds of *Health Care Activity*, such as *Diagnostic Procedures* or *Therapeutic Procedures*. Such partitions categorise clinical tasks into different kinds.

3. Contextualizing Task-Based Ontologies

3.1 Generic Context based on UMLS

We used the CPG for the Management of Hypertension prepared in 1999 by the WHO International Society of Hypertension [11] to create context-based task ontologies (CTOs), which is to say ontologies reflecting the multiple contexts in which guidelines play a role. Given tasks need not be represented in the same way in different ontologies. Thus within the CTO dictated by the UMLS Semantic Network, tasks are linked only via the *is-a* and *part-of* relations, since these are the only relations which the Network itself allows. Only those *is-a* and *part-of* relations are included which pertain to classes belonging to the three Semantic Types mentioned above, along the lines illustrated in Figure 4. The Network does not specify the task-subtask hierarchy associated with any specific guideline. Rather, it provides only the generic context for more detailed ontologies.

Each of the procedures mentioned in Figure 4 can itself serve as the basis of a more fine-grained partition based on recognizing constituent parts, as discussed in [12, 13]. Thus we can distinguish between *Diagnostic Procedures involving Forecast of Outcome* and *Diagnostic Procedures not involving Forecast of Outcome*, or between *Diagnostic Procedures involving Diabetes Mellitus* and *Diagnostic Procedures not involving Diabetes Mellitus*, and so on.

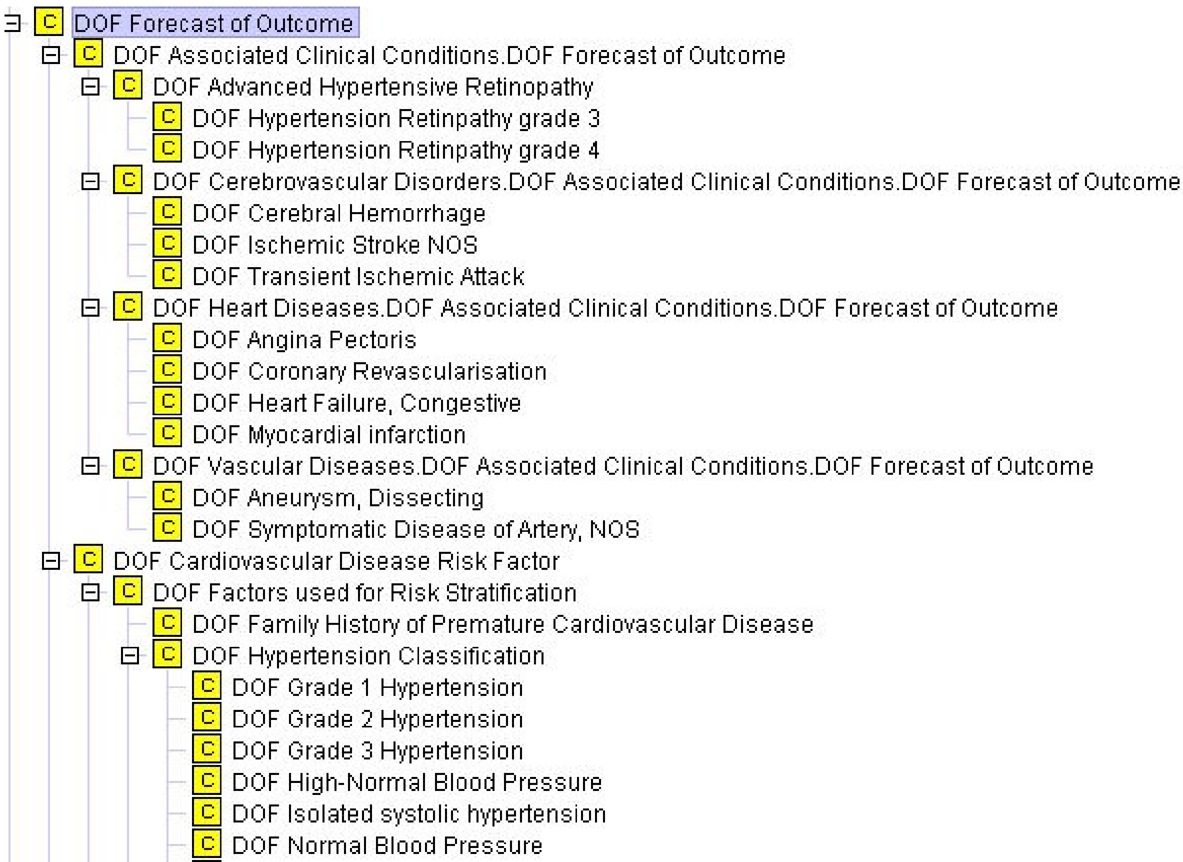

Figure 5 Context-specific task hierarchy based on WHO guideline on hypertension

3.2 Specific context based on the guideline

The task-subtask hierarchy pertinent to the WHO guideline is illustrated in Figure 5. Because we are working in the context of a hypertension guideline, the *DOF Forecast of Outcome* that is a type of *Diagnostic Procedure* has to be understood here as *DOF Forecast of Outcome of Hypertension*. We then have:

DOF Forecast of Outcome involved in Hypertension **is-a** Diagnostic Procedure

and

DOF Forecast of Outcome of Hypertension **part-of** Diagnostic Procedures involved in Hypertension

A representation of this type tells us in which context given subtasks are carried out. For instance, *Determination of Cerebral Hemorrhage* is a type of and is out carried within the context of *Determination of Cerebrovascular Disorders*, which in turn is a type of and is carried out within the context of *Diagnostic Procedure* (Figure 6).

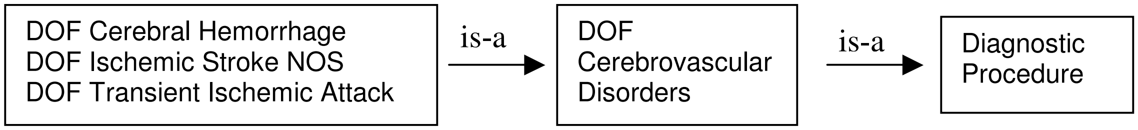

Figure 6 Guideline-specific CTO

This representation can be further extended by means of part-of relations:

> DOF Cerebral Hemorrhage **part-of** DOF Cerebrovascular Disorders involving DOF Cerebral Hemorrhage partition: DOF Cerebral Hemorrhage)
> DOF Ischemic Stroke NOS **part-of** DOF Cerebrovascular Disorders involving Ischemic Stroke NOS (partition: Ischemic Stroke NOS)
> DOF Transient Ischemic Attack **part-of** DOF Cerebrovascular Disorders involving DOF Transient Ischemic Attack (partition: Transient Ischemic Attack)

Typological relations can then be specified as follows:

> DOF Cerebrovascular Disorders involving DOF Cerebral Hemorrhage **is-a** DOF Cerebrovascular Disorders
> DOF Cerebrovascular Disorders involving DOF Ischemic Stroke NOS **is-a** DOF Cerebrovascular Disorders
> DOF Cerebrovascular Disorders involving DOF Transient Ischemic Attack **is-a** DOF Cerebrovascular Disorders

And these lead to:

> (DOF Cerebrovascular Disorders involving DOF Cerebral Hemorrhage)
> ∪ (DOF Cerebrovascular Disorders involving DOF Ischemic Stroke NOS)
> ∪ (DOF Cerebrovascular Disorders involving DOF Transient Ischemic Attack)
> **=** DOF Cerebrovascular Disorders

where '∪' signifies the operation of taking class-theoretic or taxonomic unions of classes [14].

Here the task *DOF Cerebrovascular Disorders* is represented in relation to three different sorts of subtasks, whose completion is necessary for the completion of the whole task. In this case, DOF Cerebrovascular Disorders has been classified using a series of partitions for each of which we represent the involvement and non-involvement of given subtasks. What is shown above are the relations where the subtask is involved in the task and on similar lines, there are relations where subtasks are not involved in the task. In the above cases, they include:

> DOF Cerebrovascular Disorders not involving DOF Cerebral Hemorrhage **is-a** DOF Cerebrovascular Disorders
>
> DOF Cerebrovascular Disorders not involving Ischemic Stroke NOS **is-a** DOF Cerebrovascular Disorders

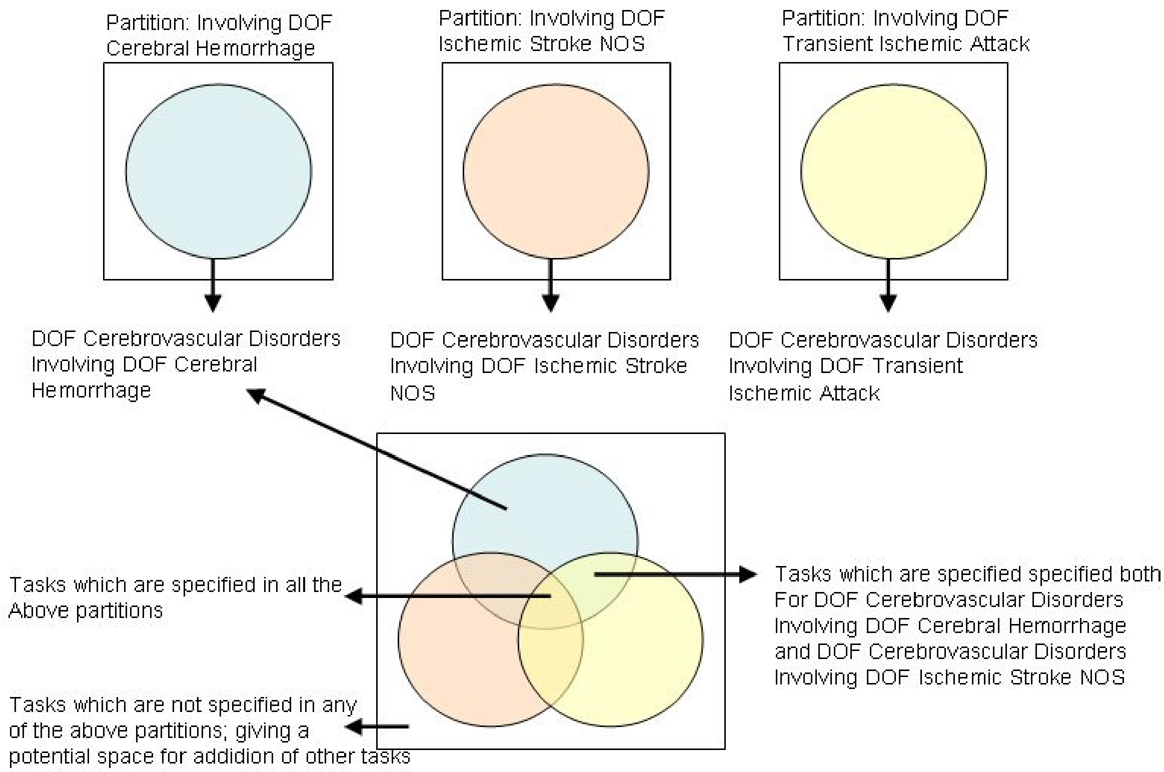

Figure 7 Partitions with Venn diagrams.

DOF Cerebrovascular Disorders not involving DOF Transient Ischemic Attack **is-a** DOF Cerebrovascular Disorders

Adding these negative clauses means that task representations can be more specific. They also allow us to assert that two classes have no class-theoretic intersection, for example as follows:

DOF Cerebrovascular Disorders involving DOF Transient Ischemic Attack $\cap$ DOF Cerebrovascular Disorders not involving DOF Transient Ischemic Attack **=** $\varnothing$ (null)

DOF Cerebrovascular Disorders involving DOF Transient Ischemic Attack $\cup$ DOF Cerebrovascular Disorders not involving DOF Transient Ischemic Attack **=** DOF Cerebrovascular Disorders

When such relations are defined within a given partition, we can exploit our machinery for representing the positive and negative involvements and unions in order to specify given tasks relative to given contexts in a very precise manner. We can also use Venn diagram representations as in Figure 7, where the rectangles in each case represent the relevant universe of discourse, the circles represent given DOF-classes, and the empty spaces represent those portions of the task space which are not involved in any subtasks considered within the respective partitions. This leaves room for the addition of new tasks, which might be needed when we move out of the context as strictly defined by the current guideline. (See the treatment of 'empty space' in [13].) Overlaps signify cases where a single subtask is carried out in the implementation of a plurality of larger tasks within a given health care setup; for example patient history taking forms part of all the three tasks considered in the above. Some of the higher-level entities comprehended by guideline representation systems such as Guide [15, 16] are modeled as *intentions* in guideline modeling standards like Asbru [17] or as *plans* in Proforma [18]. Here, however, we do not distinguish between intentions, tasks and plans, though our use of the framework of granular partitions [13] means that this and similar distinctions can easily be introduced into the theory via a new layer of refinements.

3.3 Various views to interpret and implement clinical practice guidelines

Clinical practice guidelines (CPGs) are composed always in relation to a generic clinical context. However each specific implementing health care organization will look at each given CPG from a different point of view. Thus an internal medicine team would need to supplement the CPG with details designed to align it with the workflow practices in the specific health care set-up. For example, Medical History Taking can form a part of DOF Cerebrovascular Diseases, as a result of which the following relations would need to be represented in the corresponding ontology:

> Medical History Taking involved in DOF Cerebrovascular Diseases **part-of** DOF Cerebrovascular Diseases involving Medical History Taking (Partition: Medical History Taking)
>
> DOF Cerebrovascular Diseases involved in Medical History Taking **part-of** Medical History Taking involving DOF Cerebrovascular Diseases (Partition: DOF Cerebrovascular Diseases)

where 'A involved in B' and 'B involving A' are noun phrases designating, respectively, the class of those As existing as parts of Bs, and the class of those Bs with As as parts. Similar relations apply also to the subtasks. Thus for DOF Transient Ischemic Attack, the relations include:

> Medical History Taking involved in DOF Transient Ischemic Attack **part-of** DOF Transient Ischemic Attack involving Medical History Taking (Partition: DOF Transient Ischemic Attack)
>
> DOF Transient Ischemic Attack involved in Medical History Taking **part-of** Medical History Taking involving DOF Transient Ischemic Attack (Partition: DOF Medical History Taking)

The parthood relations are maintained as before and thus:

> Medical History Taking involved in DOF Transient Ischemic Attack **part-of** DOF Cerebrovascular Disorders involving Medical History Taking (Partition: Medical History Taking)

This reveals the existence of finite chains of parthood relations revealed by the partitions at issue, and it shows how the implementation of tasks in an actual healthcare setup is linked to the semantic relationships present in the guideline. It also shows how a plurality of ontologies or classification schemes can be created reflecting the different contexts in which the ontologies need to be applied.

There can be other views: nursing staff will have a different task ontology for the hypertension guideline, which would also include slots for inpatient ward care for complicated cases, for home visits, counseling, and so on.

3.4 CTOs and clinical practice guidelines formalism

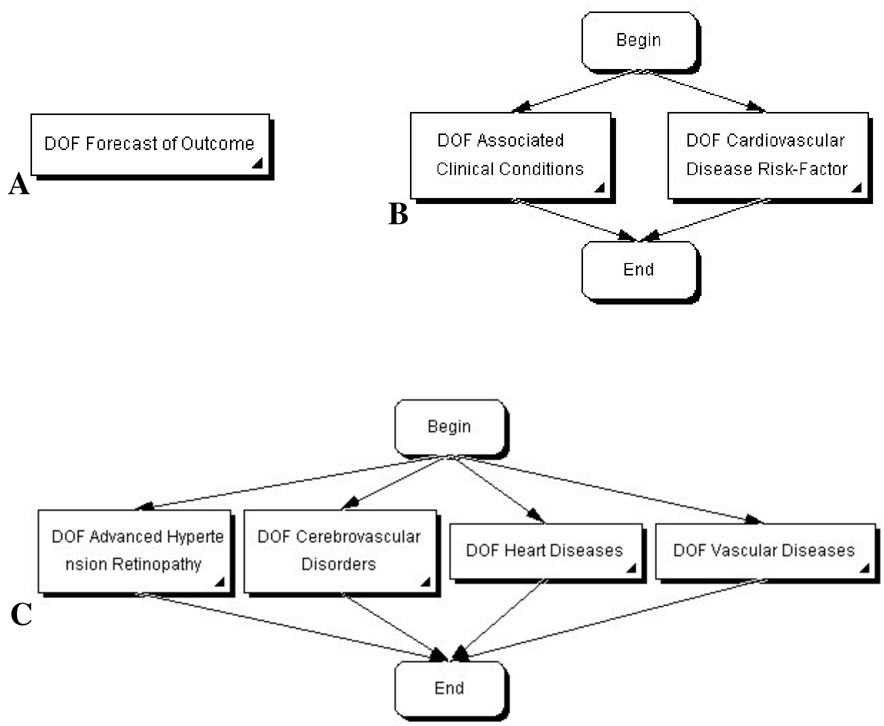

Figure 8 Tasks represented in the Guide formalism

A guideline serves as the basis for the creation of a number of CTOs that can then be used to support reasoning about tasks and subtasks which are required by those implementing the guideline from different points of view. The framework sketched above does not provide the resources for representing schedule and other details related to the particular contexts of implementation. For these purposes it will need to be supplemented by other types of machinery for translating CPG recommendation into a computerized format, for example by the Guide formalism [15, 16], which allows us to use our framework to build a flow of tasks along the lines illustrated in Figure 8. First, we define a complex task called *DOF Forecast of*

Outcome (A). The bottom-right black triangle means that the task is complex and it will be defined by a hierarchy of tasks and subtasks (in Guide each expansion is called a level). Sublevel B specifies the two complex tasks involved in the *DOF Forecast of Outcome*. C defines the complex task in the B level *DOF Associated Clinical Conditions*. All the information related to time and to the order of execution may be specified by using the Guide authoring tool. In the case of hypertension diagnosis, all the tasks can be performed in parallel as no particular order of execution or time schedule is required. Moreover the situation is quite simple, because all the DOF activities defined in the "forecast" are concentrated in a particular part of the guideline.

A different situation arises in the management of hypertension. In this case, the sequence of tasks and their conditional dependency are more important and more complex.

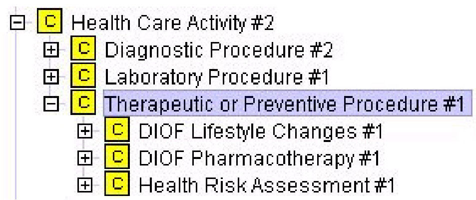

Figure 9 CTO regarding Health Care Activity related to the WHO Hypertension guideline.

Let us consider a portion of the context-specific ontology regarding the health care activities involved in hypertension management (Figure 9). This is a formal ontological representation, but for purposes of implementation in a given health care set-up we need to add information related to the sequence and scheduling of the tasks to be performed. It is clear that "Laboratory Procedure" includes a wide variety of different tasks referred to within the guideline as a whole, and it is no less clear that these are spread along the temporal axis in the management of hypertension in such a way some tasks can exist only as successors to other, already completed tasks.

For instance, in the hypertension treatment, the lifestyle changes are always needed, while the pharmacotherapy is strictly recommended only for patients in the high risk group. For all other patients the decision to perform pharmacotherapy is up to the physician. We can take into consideration at least two ways of modelling the time schedule in this context.

The left part of Figure 10 depicts a situation where lifestyle modifications are suggested in parallel with a possible pharmacotherapy. The synchronization block (triangle labelled "T") recommends the initiation of the subsequent parallel tasks at some specific point in time. The diamonds are related to decisional processes either rule-based (single line diamonds as for example in the case: *Risk Group Evaluation*) or user-based (represented by double line diamonds as in *Consider Pharmacotherapy*). In the first case the system will take a decision on the basis of specific rules. In the latter the system will ask the user what will be the next activity to perform.

The same tasks are represented also in the right part of the figure, but capturing a particular health care set-up context, where the communication of recommended lifestyle changes is usually carried out as the first task in the treatment.

While the ontology is the same in both cases, the time schedule and possibly the decision processes involved will differ.

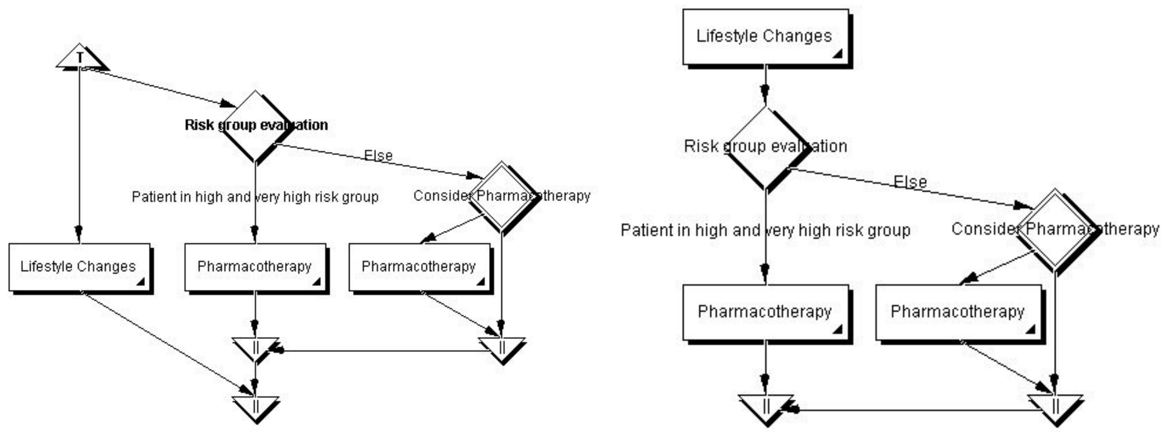

Figure 10 Two different formalizations of the same steps of the WHO Hypertension guideline

3.5 Multiple clinical practice guidelines CTOs merging

Finally a challenging topic is that of merging different CPGs together. The merging process is complex and it is related to both ontologies and time flows. In fact it is not sufficient to understand that given tasks present in two different CPGs are similar; it is important to understand also what their respective time schedules are. Clearly, successful merging of ontologies needs careful attention to temporal issues (Figures 11 and 12).

Figure 11 Example of CTO for Diabetes Mellitus guideline.

Ontologies can help us to understand the relations between tasks involved in the execution of different clinical practice guidelines. If, for example, a patient is diabetic and hypertensive, then there are at least two possible CPGs that can be taken into consideration. To manage such cases, we used the Prompt suite of tools for multiple-ontology management that is provided as an extension to the Protégé ontology-editing environment in the form of a set of plug-ins. The iPrompt algorithm takes as input two ontologies and guides the user in the creation of one merged ontology as output. The slots of the classes merged are retained as they are within the original ontologies. Using these tools, the overlapping portions of the hypertension and diabetes CPGs were mapped and further relations created within the tasks originally existing within the two separate ontologies [19].

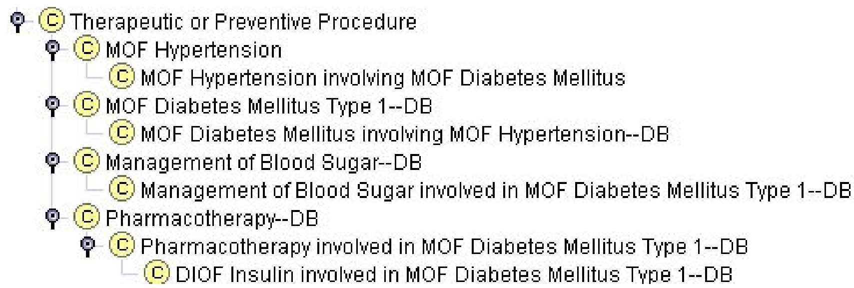

Figure 12 Diabetes Mellitus guideline CTO merged with a portion of the Hypertension guideline CTO

4 Advantages and Disadvantages of CTOs

The advantages of context-based task ontologies include:

1. The CTOs are able do justice to the fact that the same clinical task can be performed within different contexts and to the fact that the significance of a given task within those different contexts may itself be different.
2. The CTOs do not provide some unique ontology structure. Rather, they reflect the fact that there are different views shared among the different specialists involved in the implementation of clinical practice guidelines.
3. CTOs separate the specification of a task without temporal annotation and the carrying out of a task with temporal annotation. Task ontologies are thus divided into two levels, allowing the separation of the issues pertaining to the generic guideline from those pertaining to implementation within a given health care environment.
4. The CTO framework provides a robust methodology which can form the basis of integration of clinical guideline recommendations with the actual clinical practices in a given health care setup.
5. CTOs can provide a basis for merging various task ontologies from different clinical guidelines. This can be useful in the management of patients with multiple clinical disorders. We are doing further work in this direction.
6. CTOs are application independent and thus they can be used by the various tools which create Computer Interpretable Guideline Models.

On the other hand, the disadvantages of CTOs include:

1. Specification of parthood and typological relations is time-consuming and it is difficult to automate this process.
2. Creation of CTOs needs the help of experts from the relevant disciplinary fields/contexts.
3. Providing multiple ontologies is a lengthy task and the execution of computer-supervised comparisons and mergers is difficult.

5 Concluding Remarks

Context specification is difficult, especially when it comes to decision support and the implementation of specific tools within a healthcare organization [20]. Almost every guideline modelling tool specifies the context in form of pre-conditions and the rules themselves. However, these are not formalized, and communities like HL7 have many Special Interest Groups currently dealing with the problem of context specification [21].

In fact failure to specify clearly the context of application is one of the main reasons why decision support tools have not been well accepted by the medical community. In this paper, we have offered a means of specifying context based on representing clinical tasks, which form only a part of the larger clinical context. Such formal means need to be found to represent each of the other aspects relevant to clinical decision support. Recently GELLO [22] has been proposed within HL7 as a language for task specification, and efforts are underway to use it to represent clinical contexts. Irrespective of the language, one needs to find more robust ways to represent clinical contexts than those thus far developed.

References

[1] Institute of Medicine (1990). *Clinical Practice Guidelines: Directions for a New Program*, M.J. Field and K.N Lohr (eds.) Washington, DC: National Academy Press (page 38)

[2] Shea S, DuMouchel W, Bahamonde L. A Meta-analysis of 16 Randomized Controlled Trials to Evaluate Computerbased Clinical Reminder Systems for Preventative Care in the Ambulatory Setting. J Am Med Inform Assoc. 1996;3(6):399-409.

[3] UMLS website http://www.nlm.nih.gov/research/umls/

[4] DAML+OIL website http://www.daml.org/2001/03/daml+oil-index.html

[5] Kumar A, Ciccarese P, Quaglini S, Stefanelli M, Caffi E, Boiocchi L. Relating UMLS semantic types and task-based ontology to computer-interpretable clinical practice guidelines. Proceedings of Medical Informatics Europe, St.Malo, France, 4-7 May 2003, IOS Press. : 469-474,2003.

[6] OilEd website http://oiled.man.ac.uk/

[7] Bechhofer S, Horrocks I, Goble C, Robert S. OilEd: a Reason-able Ontology Editor for the Semantic Web. Proceedings of KI2001, Joint German/Austrian conference on Artificial Intelligence, September 19-21, Vienna. Springer-Verlag LNAI Vol. 2174, pp 396--408. 2001.

[8] Smith B and Mulligan K, Framework for formal ontology, Topoi, 3, 73-85, 1983.

[9] Fellbaum, C. (2002). Parallel Hierarchies in the Verb Lexicon. Proceedings of the OntoLex Workshop , LREC, Las Palmas, Spain (2002).

[10] Trautwein M and Grenon P. Roles: One Dead Armadillo on WordNet's Speedway to Ontology. GWC 2004, Brno, Czech Republic, January 20-23, 2004 (Accepted).

[11] WHO Hypertension Guideline

[12] Bittner, T. and Smith, B. (2003a). Granular Spatio-Temporal Ontologies, To appear in: Proceedings of the AAAI Spring Symposium on Foundations and Applications of Spatio-Temporal Reasoning (FASTR).

[13] Bittner, T. and Smith, B. (2003). A Theory of Granular Partitions. In: Foundations of Geographic Information Science, M. Duckham, M. F. Goodchild and M. F. Worboys (eds.), London: Taylor & Francis, 117–151.

[14] Smith B, Rosse C. The Role of Foundational Relations in the Alignment of Biomedical Ontologies, Proc Medinfo (in press).

[15] Peleg M., Tu S, Bury J, Ciccarese P, Fox J, Greenes RA, Hall R., Johnson PD, Jones N., Kumar A., Miksch S., Quaglini S., Seyfang A., Shortliffe EH, and Stefanelli M. Comparing Computer-Interpretable Guideline Models: A Case-Study Approach. J Am Med Inform Assoc. 2003 Jan-Feb;10(1):52-68.

[16] Quaglini S, Dazzi L, Gatti L, Stefanelli M, Fassino C, Tondini C. Supporting tools for guideline development and dissemination. Artif Intell Med. 1998;14(1-2):119-37.

[17] Shahar Y, Miksch S, Johnson P. The Asgaard Project: A Task-Specific Framework for the Application and Critiquing of Time-Oriented Clinical Guidelines. Artif Intell Med. 1998;14:29-51.

[18] Bury JP, Saha V, Fox J. Supporting ‘scenarios’ in the PROforma guideline modelling format. Proc AMIA Annu Symp. 2001:870.

[19] American Diabetes Association. Screening for Type 2 Diabetes. Diabetes Care 26:S21-S24, 2003

[20] Kumar A, Smith B, Pisanelli DM, Gangemi A, Stefanelli M. An ontological framework for the implementation of guidelines in health care organizations. In this volume.
[21] HL7. www.hl7.org
[22] Sordo M, Ogunyemi O, Boxwala AA, Greenes RA. GELLO: An Object-oriented Query and Expression Language for Clinical Decision Support. Proc AMIA Symp. 2003; 2003: 1012.

Acknowledgement
Work on this paper was supported by the Wolfgang Paul Program of the Alexander von Humboldt Foundation.

Contact Information: Anand Kumar, IFOMIS, University of Leipzig, Härtelstraße 16-18, 04107 Leipzig, Germany. Email: <anand.kumar@ifomis.uni-leipzig.de>