

An Ontological Framework for the Implementation of Clinical Guidelines in Health Care Organizations

ANAND KUMAR

Institute for Formal Ontology and Medical Science, University of Leipzig, Germany.
Email: anand@aim.unipv.it

BARRY SMITH

Institute for Formal Ontology and Medical Science, University of Leipzig, Germany.
Department of Philosophy, University at Buffalo, New York, USA.
Email: phismith@buffalo.edu

DOMENICO M. PISANELLI

Laboratory for Applied Ontology, Institute of Cognitive Sciences and Technologies,
National Research Council, Italy.
Email: d.pisanelli@istc.cnr.it

ALDO GANGEMI

Laboratory for Applied Ontology, Institute of Cognitive Sciences and Technologies,
National Research Council, Italy.
Email: a.gangemi@istc.cnr.it

MARIO STEFANELLI

Laboratory of Medical Informatics, Dept. of Computer Science, University of Pavia,
Italy.
Email: mstefa@aim.unipv.it

Forthcoming in D. Pisanelli (ed.), *Medical Ontologies*, Amsterdam: IOS Press.

Abstract: The paper presents the outlines of an ontology of plans and guidelines, which is then used as the basis for a framework for implementing guideline-based systems for the management of workflow in health care organizations. The framework has a number of special features, above all in that it enables us to represent in formal terms assignments of work-items both to individuals and to teams and to tailor guideline to specific contexts of application in health care organizations. It is designed also to enable implementations to do justice to the fact that the processes carried out in health care organizations may deviate in different ways from the norms set forth in corresponding guideline definitions. This means that implementations built in conformity with the framework will be marked by a type of flexibility that might make them more likely to be accepted by healthcare professionals than are standard guideline-based management systems.

1. INTRODUCTION

Guidelines for clinical practice are becoming ever more popular in every sector of health care. Guidelines have the goal of indicating the decisions and tasks most appropriate for optimizing health outcomes and for controlling costs. They can be expressed either in the form of textual recommendations or as protocols or flow diagrams. They can consist either in loose indications of a preferred set of choices or in normative rules requiring more or less strict adherence.

Clinical practice guidelines are seen by many as important tools for improving the quality and cost-efficiency of care in an increasingly complex environment of health care delivery, and it has been suggested that adherence to guidelines may reduce the costs of care by up to 25% (Clayton & Hripcsack 1995). Moreover it has been pointed out that information overload, the constraints of timeliness, and the high human and financial costs of medical error mean that it will become increasingly difficult for physicians to practice high-quality evidence-based medicine without the aid of computerized decision support systems at the point of care (Cartwright *et al* 2002).

Many efforts have been devoted in recent years to the establishment of systems for the formal representation of guidelines that will enable the implementation of computerized tools aimed at supporting clinical care and resource management on the basis of the most up-to-date medical information (Peleg *et al.* 2002). Computerization should increase the effectiveness of both the retrieval of the information contained in guidelines and the delivery of guideline-based care.

The task of matching healthcare processes to patient needs has led further to an increased interest in healthcare information systems that are easily adjustable to changes in resource availability and in organizational structures. (Poulymenopoulou *et al.*, 2003). Different European studies have shown the positive results that can be gained from using commercial workflow systems (Haux *et al.*, 2003, Lechleitner *et al.*, 2003, Muller *et al.*, 2001) and computer-interpretable workflow models based on clinical guidelines have been implemented in fields such as stroke (Quaglino *et al* 2001), post-stroke rehabilitation (Panzarasa *et al* 2002), diabetes (Barahona *et al* 2001), and radiation therapy (Karlsson & Eklund 2001). (Quaglino *et al* 2000, 2000b) have formulated careflow models in terms of Petri Nets, pointing out that guideline implementation should serve not only the issuing of simple reminders but also as an organizational template that facilitates health care processes in an environment within which responsibilities are widely shared and in which health care professionals may for a variety of reasons be non-compliant with guidelines.

In an optimal scenario guidelines would be integrated with the information systems which are in any case operational at the point of care. The full potentialities of computerized systems can then be exploited in such a way that they can cope, for example, when a range of different processes need to be executed in parallel on different patients. Such systems must be able to retrieve in real time the situation of every patient as well as give an overall report on the situation in the ward as a whole. By keeping track of the parallel tasks performed, systems of this sort would be able to prevent unnecessary duplication and warn of possible errors and omissions, thereby freeing physicians to concentrate on making clinical decisions. An optimal system, too, would be flexible in the sense that it would allow for up-dating of the guidelines themselves in tandem with advances in medical knowledge.

2. AN ONTOLOGICAL FRAMEWORK FOR CLINICAL GUIDELINES

In what follows we offer an ontological framework for the construction of computerized guideline support systems. Well-constructed ontologies serve not only to make it easier to re-use existing knowledge in new settings, but also as the foundation for standardization efforts, since they make explicit the conceptualizations behind given terminologies or models. What is required, in fact, is not some single all-powerful conceptualization of some relevant domain, but rather a reliable formal framework for the unambiguous communication of complex and detailed concepts that would work behind the scenes at the level of computer processing. The relevant information would then be conveyed to human users informally (possibly in a variety of different natural-language representations), leaving the latter free to associate their own conceptualizations according to context.

Gangemi and co-workers (1999) developed to this end ONIONS, a methodology for integrating domain terminologies by exploiting a library of generic ontological theories. Using this methodology they created ON9.2, a library of both general and domain specific ontologies, of which the web of ontologies related to guidelines sketched in Figure 1 forms part (Pisanelli et al. 2000).

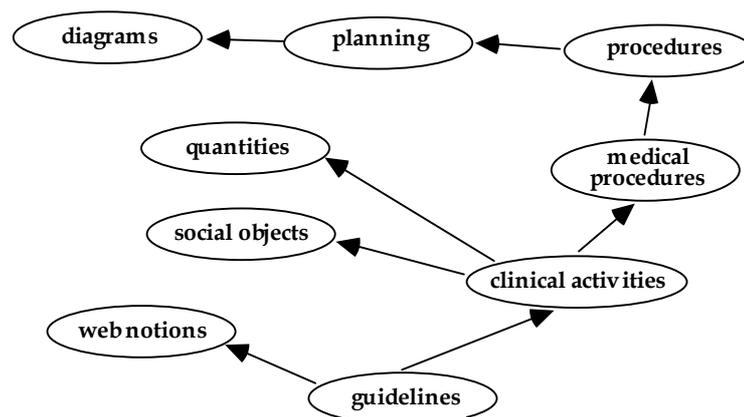


Figure 1: An extract from the ontology library ON 9.2.

Each oval represents an ontology, i.e. a module which embodies the formal definitions of related classes and relationships. Arrows denote inclusion relations between ontologies, a more specific ontology includes a more generic one.

The ontology of *medical procedures*, for example, includes that of *procedures*, which means that the classes and relationships of the latter are used by the former. The ontology of *guidelines*, similarly, includes that of *clinical activities*. The latter includes classes such as: *patient*, *patient group*, *health care operator*, *medical device*, *health care structure*, *medical sign* and *health condition*. In addition it includes relations such as: (1) *treats* and (2) *diagnoses* (between *health care operators* and *health conditions*); (3) *cares for* (between *health care operators* and *patients* or *patient groups*).

The ontology *medical procedures* comprehends also definitions of the main classes of medical procedures, drawing on the UMLS Semantic Network (NLM 2002) and on the definitions supplied by the Institute of Medicine (<http://www.iom.edu>). Five main sub-

classes of medical procedure are distinguished: (1) screening and prevention procedures; (2) surgical procedures; (3) care of clinical conditions procedures; (4) diagnostic procedures; (5) laboratory procedures.

We can understand the ontological structures of the instances of these classes in terms of the dependence relations in which they stand to entities in other classes. Thus instances of classes (1) to (3) share the need for some bearer (a patient or patient group), for some pharmacologic resource, and for some medical device. Instances of class (4) involve a patient or patient-group as target of a diagnostic action carried out on certain clinical conditions. Instances of class (5) involve certain chemical or bodily substances and a process of analysis thereof. Relations such as: *bearer*, *target*, *analyzes*, and *diagnostic action* are defined according to the more generic ontology of *agents*, which concerns the entities which play particular *functional roles* (e.g. as performer, instrument, goal, etc.) in given processes.

What is peculiar to the ON9.2 library of ontologies is its integration of a medical domain ontology within a wider framework of domain-independent foundational ontologies. The latter includes the formal resources of a mereology (or theory of part-whole relations), a topology (or theory of boundary and connectedness relations), a morphology (or qualitative theory of shapes and congruence), as well as theories of agency, planning, and so on. Recently, ON9.2 has been aligned to the DOLCE foundational ontology (Masolo et al. 2002) and efforts are underway to align it with the Basic Formal Ontology (BFO) which is being developed in Leipzig.

3. THE ONTOLOGY OF PLANS

Foundational theories have proved to be essential to the coherent integration of heterogeneous guideline models within a single platform and they seem to be necessary to the formulation of a code of best practices in the development of ontologies in general.

As an example, the Asbru model (Shahar et al. 1998) defines a number of guideline properties, but it divides them in what seems to be an unmotivated way into two sets: one (plan *status*) includes ‘rejected’, ‘ready’, etc., another (plan *state*) includes ‘aborted’, ‘completed’, etc.. Understanding such properties within an ontological framework requires at least that one recognizes the difference between *plans* and *procedures*.

A plan is an *intentional object*, which survives identically through time as a publicly available enduring template to be *realized* through the execution of a sequence of procedures in a given *situation*. The plan is represented in the physical form of written or spoken natural language descriptions, diagrams, computer applications, or in some other way.

A situation is constituted by a temporal succession of actions and/or processes actually occurring in the real world, their agents and resources, their inherent spatio-temporal regions, and other possible attributes.

A plan involves branch-points at which decisions must be taken as between alternative paths in light of prevailing circumstances; a procedure, in contrast, involves only one such path. Each successive step within a plan may be executed through a range of different actions, provided only that the constraints of the plan are satisfied. Thus different kinds of agents may be involved in the execution of a plan.

Once plan and procedure are recognized as being distinct, then one can easily see that the two sets of Asbru properties actually imply the distinction in question: only a plan can be ‘rejected’, while only a procedure can be ‘aborted’.

The above distinctions have been drawn in the context of the ON9.2 core ontology of bio-medicine (Gangemi et al. 1999, Pisanelli et al. 2000), and they are also present in the so-called DOLCE-Lite+ foundational ontology library (Gangemi et al. 03).

In DOLCE-Lite+ the class *plan* is a subclass of the class of *descriptions*, the latter being intentional objects that provide a context for formalizing several kinds of constraints on DOLCE classes (objects, substances, activities, processes, regions). A plan, for DOLCE, is composed of constraints, the latter instantiating classes such as *task* and *functional role*. The relation between a description and a situation is called *satisfies*, and in the context of planning can be specialized as an *executed according to* relation (Figure 2). The *planning* ontology contains also a number of further distinctions, for example between *simple* and *complex* tasks, and between real world and reasoning tasks, the latter also in accordance with the Pen&Pad model for patient records of (Rector et al. 1990).

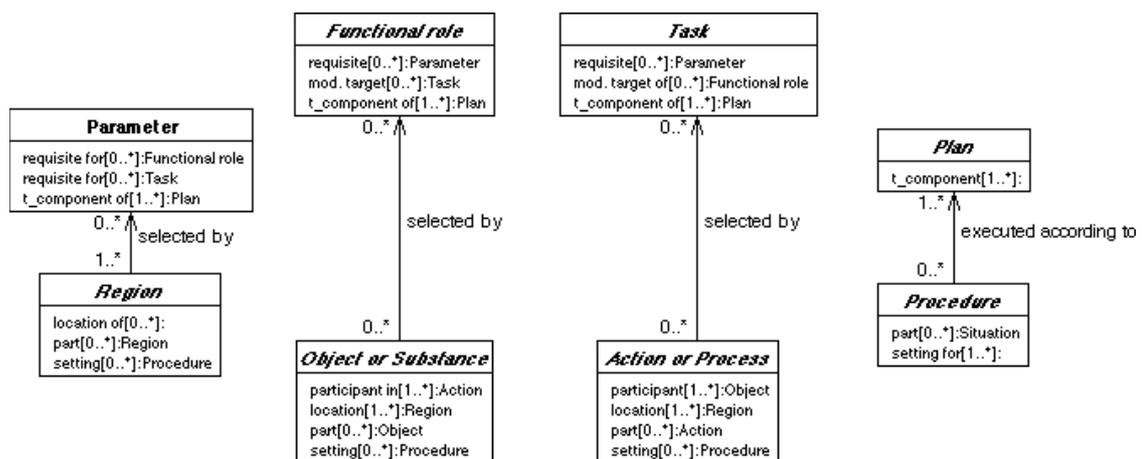


Figure 2: A UML class diagram for the plan ontology.

A plan contains constraints to a procedure components (its “setting”). A task constrains (“selects”) the execution or verification of actions/processes; a functional role constrains the participation of objects or substances; a parameter selects a certain range of values in region.

4. TEAM-ENABLED WORKFLOW SYSTEMS

Thus (Stefanelli 2001, 2002) has argued that a guideline should be viewable within the context of a specific health care organization in such a way as to yield a system for management of work activities across the organization.

The present paper represents the first step in the attempt to expand the ONIONS/DOLCE project in such a way as to include a series of further concepts specific to the domain of workflow systems, and also further domain-independent elements deriving from the BFO ontology, and in particular from its theory of granular partitions discussed below.

Several relations can be defined among different plans executed on similar or overlapping sets of entities. In the present paper, we will mainly demonstrate how one particular such relation – the relation of *refinement* – is applied in the domain of workflow systems, a typical area of guideline exploitation.

When plans are applied in specific real-world contexts then they become refined. That is, we move from coarse- to fine-grained descriptions as we raise the degree of detail and specificity in our representation of the components of the plan, for example by setting more constraints on the constituents (including agents) involved in the execution of the corresponding procedures.

A guideline, now, is a certain kind of plan. And like plans in general so also guidelines in particular are subject to the process of refinement, beginning with the guideline description as it emanates from a body such as the National Institutes of Health, and concluding with the sequences of specific work-orders prepared on site from one day to the next. Guidelines must accommodate in this process the factor of *flexibility*. In the context of our present discussion this means techniques must be developed which will allow the integration of guideline-based recommendations into the management of health-care processes in ways that are sensitive to the needs of specific patients and specific health care organizations (Sim *et al.*, 2001; Wald *et al.*, 2001).

Workflow systems typically employ a tripartite categorization of: *cases*, *work-items* and *resources*. A *case* is a specific situation in which the workflow system is applied; a *work-item* is a task to be performed in relation to a given case; *resources* are the persons and facilities needed to execute given work-items.

To understand these classes in relation to the ontology framework outlined above, we note first of all that while *plans* and *guidelines* belong to the level of descriptions: they are abstract descriptive contents. *Cases*, *work-items*, and *resources*, in contrast, are concrete real-world entities. Provisionally, we can think of a plan as what is represented in a description of a branching structure of the type mentioned above involving *tasks* to be performed by certain *agents* having certain *functional roles*. When a plan is *applied* in a given health care situation, then the tasks are transformed into work-items, in the description of which certain specifications left open in the original task description are concretized to the specific case in hand. When this work-item description is *executed* this gives rise to a temporal sequence of actions and processes.

As concerns the dimension of *functional roles*, and therewith also that of *functions*, we note that not every agent within a health care organization can carry out every task. Agents have different functional roles for example according to the sorts of tasks which they are authorized to carry out. Their ability to carry out a range of tasks of given sorts in a given situation is called their *function*.

Existing applications standardly assign work-items not to teams but to specific workers. An application may in some way recognize that teams exist, but it is pre-selected single members of teams who are called upon to execute specific work-items at specific times. In reality, however, team-work – and the flexibility that goes together therewith – is one of the key characteristics to be exploited by a workflow system. The teams within a given health care organization have collective functions and they may be collectively responsible for their execution. To put it simply: doctors, nurses, technicians and assistants work in tandem, and current workflow models do not do justice to this fact.

In (van der Aalst and Kumar, 2001) a reference model for team-enabled workflow systems is proposed in which a team is conceived as a collection of resources with

collective functional roles. They provide a concept of *team_type*, which refers to a structure which corresponds to a collection of resources. Each work-item is then associated not with an individual agent but with an instance of *team_type*. For example, a group of internists, nurses and nurse assistants in the outpatient ward of a tertiary hospital is an instance of *team_type*, and so is the surgical team in the operating theater.

In what follows we propose a general framework for guideline-based health care organization management within which it will be possible to understand in formal terms what it is to assign work-items to teams and to individuals. In addition the framework is designed to do justice to the fact that the actual realizations of workflow processes may deviate from the norms set forth in the process definitions themselves. The implementation is stratified, so that it lets us define different levels at which clinical guidelines are described in the course of implementation within a given organization.

5. PARTITIONS, VAGUENESS AND APPROXIMATION

Different partitions at different levels of refinement are needed in order to highlight different aspects of a health care organization and of its workflow processes. At the same time a framework is needed within which these different partitions can be manipulated simultaneously. To this end we employ the Theory of Granular Partitions (TGP) put forward in (Bittner & Smith 2003, 2003b, 2003c, Bittner, in press). A partition, from the perspective of TGP, consists of a network of cells and subcells, the latter being nested within the former; the cells, in turn, are projected onto entities in reality or onto the cells of other partitions. Thus, the theory of granular partitions has two parts: a theory which relates cells to the larger partition structures in which they are included, and a theory which projects cells and subcells onto other entities.

Partitions are called ‘granular’ because they project onto reality in every case only at a certain level of granularity: the precise chemical structure of a prescribed pharmacologic substance, or the precise steps which need to be taken in administering this substance, are factors which characteristically fall below the threshold of salience of a given clinical guideline or work-order, and are thus traced over in the corresponding granular partitions. Plans in general are in every case granular in just this sense. (Indeed – though we do not pursue this matter here – plans themselves can be considered as having the structure of branching sequences of granular partitions, which are projected onto reality through the actions by which they are realized.)

The hierarchy of available human resources, the functions they perform, as well as the physical facilities at the disposal of an organization – all of these determine cross-cutting partitions of the health care organization which are needed for a complete ontology of team-based workflow. When a particular human resource, for example nurse A, is entitled to carry out a particular function F, then the cell labelled nurse A in the corresponding partition of human resources is projected onto function F in the partition of responsibilities. When we assess how A exercises this function, then we have a new partition, where the cell labelled nurse A is projected onto the processes she actually does perform.

The functions represented in a guideline-based workflow model are of course not always instantiated fully accurately in actual behavior. This factor is complicated in the presence of team-based workflow organization by the fact that functions may be assigned not individually but collectively. If the benefits of genuine team-work are to be realized,

then such assignments must be flexible. TGP provides the machinery to handle this problem by conceiving the relation between the abstract specifications contained in guidelines and the functions exercised by actual healthcare teams in terms of a theory of what we shall call *reference partitions*.

When we assign a function F to a team T with members A, B, C, then it is not *ex ante* clear whether F will be performed by A or B or C alone, by A in collaboration with B, by B in collaboration with C, and so on. All such possibilities might need to be left open in a team-sensitive framework for task-assignment. Alternatively it might need to be recorded that C is excluded from performing the given function except under the supervision of A or B.

Some partitions of the health care organization – for example the partitions of its physical plant or of its individual human resources – are not subject to such indeterminacy. To understand the relations between determinate and indeterminate partitions consider the way the determinate partition defined by the borders of the fifty states of the USA is used in specifying the location of an area of high pressure in a weather forecast. The boundaries of the separate states are well known and well defined, and we can use this fact to specify the indeterminate area even though we do not know exactly where it is, for example, by asserting that it overlaps fully with Texas and Arizona, partially with New Mexico, but not at all with any other state. Thus we can distinguish three sorts of specifications: of full overlap (*fo*), partial overlap (*po*) and no overlap (*no*), in terms of which the cell corresponding to a given weather phenomenon in a US meteorological partition can be projected onto a map of the fifty US States.

The same idea can be used in giving a formal account of the ways responsibilities are assigned to the members of a team within a health care organization. Here, the *ex ante* boundaries of the functions to be associated with any given member of the team are indeterminate. The *ex ante* boundaries of the actual healthcare processes which will become associated with the functions mentioned in a clinical practice guideline in a given realization are also indeterminate. But the complete list of responsibilities in the organization and the complete list of the functions determined by the guideline text are crisp, and so they can be used to determine reference partitions, analogous to the map of the 50 US States, which can be used to specify the functions of a given team of human resources in the organization in a way that still does justice to the flexibility of team-based organization.

Consider the assignment of responsibilities of nurse A within her team with regard to a given case on a given occasion. For each given healthcare function F we can distinguish three alternatives: A is required to perform the function F, A is allowed to perform the function F, and A is excluded from performing the function F. In partition-theoretic terms these three alternatives can be described as follows. The relation of A to the space of functions is analogous to the relation of the area of high pressure to the reference-map of the USA. When A is assigned the function then A's assignment projects onto the cell F with full overlap (*fo*); when A is allowed to perform the function it projects onto F with partial overlap (*po*); when A is excluded from performing the function it projects onto F with no overlap (*no*). (See (Bittner & Smith 2003c) for the formal treatment of these matters in TGP, and Bittner and Stell (2003) for the relation to the theory of Stratified Rough Sets.)

This methodology can be extended to deal with more complicated cases, for example where F is a collective function which needs to be carried out by nurse A together with other (more or less specifically determined) members of her team.

6. GUIDELINE-GENERATED PARTITIONS

Each guideline implementation presupposes a number of partitions of the corresponding health care organization:

1. *of its physical structures,*
2. *of its human resources,*
3. *of the tasks capable of being performed by these human resources,*
4. *of the tasks recommended in the guideline itself.*

Such partitions are organized in modular fashion. Thus the physical structure will standardly be divided into different departments:

of internal medicine, of surgery, of cardiology, and so on.

The department of internal medicine may in turn be subdivided into:

outpatient wards, procedure room, inpatient wards, intensive care unit, etc.

The healthcare *teams* in the organization might consist of:

internal medicine team A, internal medicine team B, general surgery team, cardiology team, etc.

Internal medicine team A might itself consist of:

physician, resident C, resident D, nurse staff E, nurse staff F, nursing student G, etc.

Similarly, the tasks performed by human resources in health care organization can be divided into:

diagnostic procedures, therapeutic procedures, etc.

Diagnostic procedure can be further subdivided into:

medical history taking, physical examination, laboratory procedures, etc.

The NIH hypertension guideline for its part divides the tasks of hypertension management into:

blood pressure measurement, classification of blood pressure, cardiovascular disease risk determination, advice: benefits of lowering blood pressure, ambulatory blood pressure monitoring, self-measurement of blood pressure, patient evaluation, treatment, management of special situations with hypertension, improving hypertension control and community program management,

and the item *cardiovascular disease risk determination* can itself be further subdivided into:

classification of blood pressure, determination of major risk factors and determination of target organ damage.

Such modular organization is captured in TGP in terms of its theory of partitions as structures of cells and subcells. Thus for example a partition reflecting a more detailed division of relevant tasks is a refinement of a partition in which only the corresponding higher-level modules are represented.

7. IMPLEMENTATION FORMALISM

The respective roots of the four trees corresponding to the four partitions mentioned above are:

1. *hospital (physical plant)*,
2. *human agents*,
3. *tasks of human agents*,
4. *tasks mentioned in the guideline*.

Based on these partitions, we can create a team-enabled workflow model relative to the NIH Hypertension Guideline along the following lines, moving from coarsest to finest level of granularity. First, we define the collective roles of the team. We consider for this purpose the *internal medicine team A*, which consists of *physician, resident A, resident B, nursing staff A, nursing staff B, nursing student A* and *nursing student B*.

G1: At the coarsest level of granularity we do not consider the interior structure of the team at all. Rather, the system records merely the fact that the given team is able to perform the functions mentioned in the given guideline within the given organisational set-up. The partitions at this level are such that a range of different sorts of detail is traced over. Thus for example it is left open whether *Medical History Taking* means the taking of a quick history of symptoms related to hypertension or an entire personal history including sleeping and drinking habits, past history of other pathologies, history of medications, family history, history related to socio-economic factors, and so on. Similarly, in regard to guideline recommendations such as: *perform cardiovascular physical examination*, it is not specified at this level what kinds of examinations (for example *inspection, palpation, percussion, auscultation*) need to be performed in any given case. It is not specified either whether these need to be carried out on each visit or on every third visit or when certain symptoms arise. The NIH Hypertension guideline does not provide precise specifications as concerns these matters, leaving an element of flexibility which has been left to the decision of health-care practitioners and teams. Some guidelines, for instance those pertaining to prenatal care, might involve more precision than the NIH hypertension guideline. Thus they might specify precisely which physical examinations need to be performed. Even in such cases, however, no maximum limit of such examinations is specified, and even the minimum limit might be specified only imprecisely.

G2: Moving to the next level of granularity in partitions of NIH hypertension guideline functions, we create a list of the functions and subfunctions related to hypertension management. Thus we recognize certain functions present in the guideline as *parts* of the functions distinguished in a coarse-grained fashion in the partitions mentioned above. *Cardiovascular disease risk determination*, for example, is recognized as consisting of:

determination of hypertension, determination of cigarette smoking, determination of obesity, determination of physical inactivity, determination of dyslipidemia, determination of diabetes mellitus, determination of microalbuminuria, determination of age and determination of family history of premature cardiovascular disease.

At this level of granularity, the functions are more precisely specified. Thus it becomes clear that *blood pressure measurement* overlaps only partially with *cardiovascular disease risk determination*, since both include *determination of hypertension* but parts of the latter, for instance *determination of obesity* or *determination of diabetes mellitus*, do not overlap with the former. We find at this level also however new dimensions of indeterminacy. Thus there is no distinct boundary between *treatment* and *management of special situations with hypertension*, or between *treatment* and *improving hypertension control*, or between *management of special situations with hypertension* and *improving hypertension control*. The function *blood pressure measurement* can designate different things according to whether it is seen as being a subfunction of the complex functions *cardiovascular disease risk determination*, *patient evaluation*, *treatment*, *management of special situations with hypertension* or *improving hypertension control*.

G3: Here we take into account also the physical structure of the hospital, the framework assigning to *internal medicine team A* in relation to a given case not merely specific functions, but also specific physical locations in which these functions are to be realized. As mentioned above, the organization in the given case consists of different departments, which for reasons of simplicity have been assumed to be mutually exclusive. Thus it is assumed that the *departments of internal medicine, of surgery, of cardiology, etc.*, together with the *central radiology center, central pathology laboratory, billing section, indents department, etc.* do not overlap. We can further take account in our partitions of the constituent parts of these departments. For example, the *department of internal medicine* might consist of:

inpatient ward 1, inpatient ward 2, procedure room, outpatient ward, and so on.

The boundaries between the physical structures are determinate; that is, one can easily demarcate where the *inpatient ward 1* ends and where the *outpatient ward* begins; it is determinate also that both of these are parts of the larger physical structure called the *department of internal medicine*. In an uncomplicated outpatient case, functions like *blood pressure measurement* would be carried out in the *outpatient ward* and so also would *determination of cigarette smoking, determination of obesity*, and so on. On the other hand, the function *determination of dyslipidemia* might contain parts such as:

medical history taking for determination of dyslipidemia,

physical examination for determination of dyslipidemia,

blood collection for determination of dyslipidemia,

advice: determination of blood lipid profile for determination of dyslipidemia,

and while the first two of these functions might be performed in the *outpatient ward*, the last two need to take place in the *central pathology laboratory*.

G4: The organization of the human resources present in the health care organization has a modular structure, being divided into submodules under the headings: *physicians, nurses, technicians, laboratory attendants, nurse students* and so on. To create workflow

representations which do justice to the existence of *teams* of human resources we need to find a way to assign functions to individuals in such a way that the flexibility characteristic of team-work is guaranteed. This does not mean that individualized roles do not exist; rather they should be seen as particular cases of team-based assignments. When functions are assigned to the members of a team, this is done in such a way as to allow for the possibilities of partial or full overlap along the lines set forth in our discussion of TGP above. *Blood pressure measurement*, for example, is a function which can be carried out by each of the members of *internal medicine team A* acting alone. This will accordingly be a function that is assigned to each member with full overlap. When we move to finer granularities, however, then different conditions apply. Thus while *blood pressure measurement* in the *outpatient ward* for an *uncomplicated case of hypertension* can be carried out by all team members, *blood pressure measurement* in *inpatient ward 1* in a *hypertensive patient with cardiac failure* could not be carried out by *nursing student A* or *nursing student B*, but only by *physician, resident A, resident B, nursing staff A* or *nursing staff B*.

Certain procedures involve more than such allocations of functions to individuals; they involve the cooperation between human beings in the performance of a single function. Such participation might be mentioned explicitly in the guideline implementation or it might be left implicit as something that is taken for granted in the given health care organization. This presence of implicit as well as explicit components can be illustrated for example in the case of a routine referral as follows:

Case referral by *resident A* to *nursing staff B* for *blood pressure measurement* (either implicit or explicit),

Blood pressure measurement by *nursing staff B* (explicit),

Case referral by *nursing staff B* to *resident A* after *blood pressure measurement* (either implicit or explicit),

Reporting by *nursing staff B* to *resident A* of *finding of blood pressure measurement* (explicit),

Interpretation by *resident A* of *finding of blood pressure measurement* (explicit),

Monitoring by *physician* of *blood pressure measurement* (implicit).

The last item is implicit because, while the physician is responsible for any mishaps during the performance of functions in the outpatient ward, this performance is not directly monitored by him and such direct monitoring is also not required in the guideline.

8. CONCLUSION

We have discussed a workflow system which involves the specification of an ordered sequence of successively more refined modular structures (granular partitions) which support the team-based assignment of the functions specified in guidelines and the tracking of the performance of these functions by human resources in a given health care organization. Some of these partitions constitute reference partitions in relation to which the specifications carried by other partitions can be defined in such a way as to accommodate different sorts of flexibility – above all the sort of flexible assignment of

functions that goes along with team-based management. A framework with such built-in flexibility is indispensable, given the ways in which the functions recommended in guideline texts are specified and assigned, if we are to build workflow systems which will be accepted by the medical community.

Some further advantages of our framework are that it enables us to do justice to the fact that the same clinical task can be performed within different contexts in such a way that the significance of the task might be different in each (for example the task of taking blood pressure measurements has a different significance in the context of an examination for diabetes examination and in the context of hypertension management). This addresses the true clinical situation in a more adequate way than competing frameworks. The approach is further embedded within a framework of foundational ontologies, so that we can use all the tools of mereology and other formal disciplines in giving an account of the different compartments of medical reality and of the wholes in which they are related together within a single, complex, dynamic whole. It provides also a framework for merging different task ontologies deriving from different clinical guidelines, which can be useful for example in management of patients with multiple clinical disorders.

The disadvantages of the framework are that it is complex, and the specification of parthood and topological and other relations is time-consuming and needs to be done manually (though our experience tells us that domain experts who are called upon to perform such texts very quickly become proficient at providing the needed information in the appropriate form). Finally (and we are not clear whether this is an advantage or a disadvantage) the framework is open-ended, in the sense that there are no limits to the partitions which can be created within its terms, and thus no point at which one can say that the ontologies are complete in themselves. One can always add new levels of detail. We can, however, employ a pragmatic notion of completeness, which is achieved when implementations built in terms of the framework do indeed perform successfully.

REFERENCES

Barahona P, Azevedo F, Veloso M, Estevao N, Gallego R. Computerising a guideline for the management of diabetes. *Int J Med Inf.* 2001 Dec;64(2-3):275-84

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

Bittner T and Smith B. Granular Spatio-Temporal Ontologies. *Proceedings of the AAAI Spring Symposium on Foundations and Applications of Spatio-Temporal Reasoning (FASTR)*, 2003b.

Bittner T and Smith B. Vague Reference and Approximating Judgements, *Spatial Cognition and Computation* (forthcoming). 2003c.

Bittner T. A Mereological Theory of Frames of Reference, *International Journal on Artificial Intelligence Tools*, in press.

Bittner T, Stell JG. Stratified Rough Sets and Vagueness. COSIT 2003.

Cartwright J, de Sylva S, Glasgow M, Rivard R, Whiting J. Inaccessible information is useless information: addressing the knowledge gap. *J Med Pract Manage.* 2002 Jul-Aug;18(1):36-41.

Clayton PD, Hripcsak G. Decision support in healthcare. *Int. J. of Bio-Medical Computing,* 1995; 39: 59-66.

Gangemi A, Mika P. Understanding the Semantic Web through Descriptions and Situations. *Proceedings of ODBASE Conference 2003.*

Gangemi A, Pisanelli DM, Steve G. An Overview of the ONIONS project: Applying Ontologies to the Integration of Medical Terminologies. *Data and Knowledge Engineering,* vol.31, 1999.

Haux R, Seggewies C, Baldauf-Sobez W, Kullmann P, Reichert H, Luedecke L, Seibold H. Soarian--workflow management applied for health care. *Methods Inf Med.* 2003;42(1):25-36.

Karlsson J, Eklund P. Workflow design as a basis for component interaction. *Medinfo.* 2001;10(Pt 2):1158-60.

Lechleitner G, Pfeiffer KP, Wilhelmy I, Ball M. Cerner Millennium: the Innsbruck experience. *Methods Inf Med.* 2003;42(1):8-15.

Masolo C, Borgo S, Gangemi A, Guarino N, Oltramari A, Schneider L. The WonderWeb Library of Foundational Ontologies and the DOLCE Ontology. *Wonderweb Deliverable D17,* available at www.loa-cnr.it, 2002

Muller ML, Ganslandt T, Eich HP, Lang K, Ohmann C, Prokosch HU. Towards integration of clinical decision support in commercial hospital information systems using distributed, reusable software and knowledge components. *Int J Med Inf.* 2001 Dec;64(2-3):369-77.

National Library of Medicine, UMLS Knowledge Sources, 2002 edition, available from the NLM, Bethesda, Maryland.

Panzarasa S, Madde S, Quaglini S, Pistarini C, Stefanelli M. Evidence-based careflow management systems: the case of post-stroke rehabilitation. *J Biomed Inform.* 2002 Apr;35(2):123-39.

Pawlak Z. Rough sets. *International Journal of Computer and Information Sciences* 1982 (11):341-356.

Peleg M et al. Comparing Computer-Interpretable Guideline Models: A Case-Study Approach, University of Stanford Medical Informatics Technical Report SMI-2002-0935, 2002.

Pisanelli DM, Gangemi A, Steve G. The role of ontologies for an effective and unambiguous dissemination of clinical guidelines. In R. Dieng and O. Corby (eds.), Knowledge Engineering and Knowledge Management. Springer Verlag, 2000.

Poulymenopoulou M, Malamateniou F, Vassilacopoulos G. Specifying workflow process requirements for an emergency medical service. J Med Syst. 2003 Aug;27(4):325-35.

Quaglini S, Stefanelli M, Cavallini A, Micieli G, Fassino C, Mossa C. Guideline-based careflow systems. Artif Intell Med. 2000 Aug;20(1):5-22.

Quaglini S, Caffi E, Cavallini A, Micieli G, Stefanelli M. Simulation of a stroke unit careflow. Medinfo. 2001a;10(Pt 2):1190-1.

Quaglini S, Stefanelli M, Lanzola G, Caporusso V, Panzarasa S. Flexible guideline-based patient careflow systems. Artif Intell Med. 2001b Apr;22(1):65-80.

Rector AL, Nowlan WA and Kay S. Unifying Medical Information using an Architecture Based on Descriptions. in RA Miller (ed) Proceedings of the 14th Annual Symposium on Computer Applications in Medical Care , SCAMC 90, IEEE Computer Society Press, Los Alamitos, California ,1990, pp 190-194.

Shahar Y, Miksch S, Johnson P. The Asgaard project, Artificial Intelligence in Medicine, 1998; 14.

Smith B, Rosse C. The Role of Foundational Relations in the Alignment of Biomedical Ontologies, <http://ontology.buffalo.edu/medo/isa.pdf>.

Stefanelli M. The socio-organizational age of artificial intelligence in medicine. Artif Intell Med. 2001 Aug;23(1):25-47.

Stefanelli M. The role of methodologies to improve efficiency and effectiveness of care delivery processes for the year 2013. Int J Med Inf. 2002 Nov 20;66(1-3):39-44.

van der Aalst WMP & Kumar A. Team-Enabled Workflow Management Systems. Data and Knowledge Engineering, 38(3):335-363, 2001.

Wald JS, Pedraza LA, Reilly CA, Murphy ME, Kuperman GJ. Requirements development for a patient computing system. Proc AMIA Symp. 2001;:731-5.

Acknowledgment: Work on this paper was carried out under the auspices of the Wolfgang Paul Programme of the Alexander von Humboldt Foundation.