Clinical Guidelines as Plans: An Ontological Theory

ANAND KUMAR
Institute for Formal Ontology and Medical Science, Saarland University, Saarbrücken, Germany.
Email: akumar@ifomis.uni-saarland.de

BARRY SMITH
Institute for Formal Ontology and Medical Science, Saarland University, Saarbrücken 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, Department of Computer Science, University of Pavia, Italy.
Email: mstefa@aim.unipv.it

Abstract

Clinical guidelines are special types of plans realized by collective agents in health care organizations. We use this idea as the basis for a new understanding of the way in which guideline-based systems can be employed in the management of workflow in health care organizations. The framework we propose allows us to represent in formal terms how guidelines are realized through the actions of individuals organized into teams. Implementations built in conformity with our framework will accordingly be marked by two dimensions of flexibility that are designed to make them more likely to be accepted by health care professionals than are standard guideline-based management systems.

1. INTRODUCTION

As has too often been pointed out, 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).

Guidelines for clinical practice will surely play an essential role in this development and such guidelines 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 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.

Much effort has 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 in a way which simultaneously takes account of the most up-to-date medical information (Peleg *et al.*, 2002). Computerization should hereby increase the effectiveness of both the retrieval of the knowledge contained in guidelines and the delivery of guideline-based care.

The task of matching health care processes to patient needs has led further to an increased interest in health care information systems that are easily adjustable to changes in resource availability and in organizational structures (Poulymenopoulou *et al.*, 2003). Different 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 (Quaglini *et al.*, 2001), post-stroke rehabilitation (Panzarasa *et al.*, 2002), diabetes (Barahona *et al.*, 2001), and radiation therapy (Karlsson & Eklund 2001).

In an optimal scenario, guidelines would be integrated with the information systems operational in a given hospital in such a way as to be made available in the most suitable form 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, in those cases where 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. They would be able to optimize the allocation of resources, generate work-schedules for specific individuals and groups in the organization – and all of this in such a way as to leave physicians free to concentrate on making the most challenging clinical decisions.

2. AN ONTOLOGICAL FRAMEWORK FOR CLINICAL GUIDELINES

Most important from our persepective is the work of Quaglini and co-workers, who have formulated careflow models in terms of Petri Nets designed to do justice to the fact that a framework for guideline implementation should serve not only the issuing of simple reminders but also as an organizational support to the management of health care processes in an environment within which (1) responsibilities are widely shared and (2) health care professionals may be non-compliant with guidelines for a variety of well justified reasons.

In what follows we present an ontological framework for the construction of computerized guideline support systems against this background. Well-constructed ontologies serve not only to make easier the re-use of already formulated knowledge in new settings, but also as the foundation for standardization efforts, since they make explicit the tacit assumptions underlying existing medical task flows. What is required, in fact, is not some single all-powerful knowledge representation formalism for given medical domains, but rather a reliable formal framework for the unambiguous representation of the complex concepts involved in models of guidelines of a sort that can support the design and development of the computational implentation of such models. The relevant information could then be conveyed to human beings informally, for example in natural-language or graphical formats tailored to the needs of each separate group of users.

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 ON9.3, a library consisting of both domain-independent and domain-specific ontologies, of which the ontology related to guidelines sketched in Figure 1 forms part (Pisanelli *et al.*, 2000).

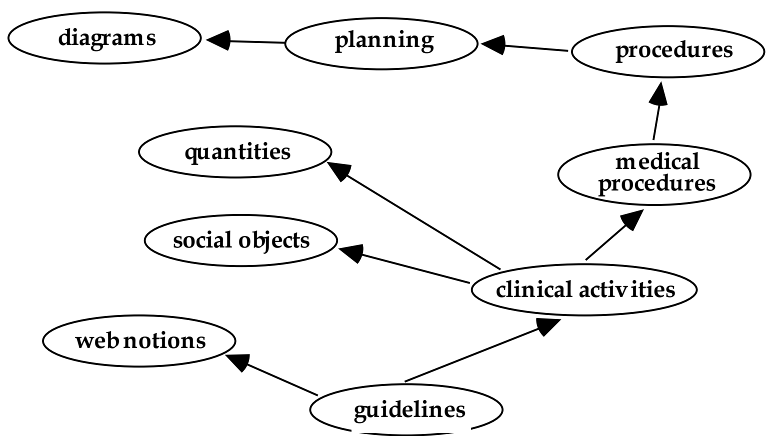

Figure 1: An extract from the ontology library ON 9.3.

Here each oval represents an ontology, i.e. a module which embodies the formal definitions of the corresponding classes and relations. Arrows denote inclusion relations between ontologies, a more specific ontology using the classes and relations taken over from a more generic one. Thus for example the ontology of *clinical activities* includes the ontology of *medical procedures*, and the ontology of *guidelines* 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:

diagnoses (between *health care operators* and *health conditions*);
cares for (between *health care operators* and *patients*).

The ontology *medical procedures* comprehends also definitions of the main classes of medical procedures, including for example *screening procedure*, *surgical procedure*, *diagnostic procedure*, *laboratory procedure*.

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 *surgical procedures* share the need for (are ontologically dependent on) instances of *surgeon*, *support staff*, and *patient*, as also of *pharmacological resources*, *medical device*, etc. Instances of *diagnostic procedure* involve a *patient* as target and a *diagnostic action* carried out on certain *clinical conditions*. Instances of *laboratory procedure* involve certain *chemical substances* and a process of *analysis* thereof.

Recently, ON9.3 has been aligned to the DOLCE foundational ontology (Masolo *et al*., 2002, Gangemi *et al*., 2003) and efforts are underway to align it with the Basic Formal Ontology (BFO) which is being developed by the Institute for Formal Ontology and Medical Information Science in Saarbrücken.

3. THE ONTOLOGY OF PLANS

What is peculiar to the ON9.3 ontology library is its integration of a medical domain ontology within a wider framework of domain-independent foundational ontologies. We believe that the development of such foundational ontologies will in due course be recognized as an indispensable element of best practices in the development of ontologies in general, and that, more specifically, they will prove to be essential for the coherent integration of heterogeneous guideline models within a single platform.

ON9.3 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), as well as the ontologies of agency and planning which will be directly relevant to us here. Classes such as: *bearer*, *target*, and *diagnostic action* are defined within the framework of the generic ontology of *agency*, which concerns the entities which play particular *functional roles* in given processes (e.g. as performer, instrument, goal, etc.), and the ontology of *plans* is then directly related thereto.

A plan is an *intentional object* with a given goal. The plan survives identically through time until it is executed in a given procedure. The plan is represented in the physical form of written or spoken natural language descriptions, web pages, and so forth, which we here group together under the heading *diagram*. If we examine the topmost part of Figure 1, then we notice that there is a dependency between the ontologies of diagrams and planning and procedures: one needs to take into account the classes of the diagram ontology in constructing the planning ontology, and the latter is needed in turn by the procedures ontology. (Gangemi *et al*., 1999, Pisanelli *et al*., 2000)

A plan involves branch-points at which decisions must be taken as between alternative paths taken by the executors of the plan in light of given circumstances; a procedure, in contrast, involves only one single 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 a multiplicity of different kinds of agents may be involved in the execution of a single plan.

Plans and procedures are clearly two different kinds of entities. Plans are something abstract. Procedures correspond to the concrete series of actions which are the executions of plans. Some guideline models in the literature have ignored this distinction. It remains as an underlying presence, however, in the Asbru framework in the distinction between what can be 'rejected' (a *plan* in our terminology) and what can be 'aborted' (a *procedure*).

4. TEAM-ENABLED WORKFLOW SYSTEMS

Stefanelli (2001, 2002) has argued that tools are required to enable a guideline to be viewed or accessed within the context of a specific health care organization in such a way as to yield a system for management of work activities within that organization at a given time.

The present paper represents the first step, against the background of the remarks on plans and procedures in the above, towards meeting this need.

A guideline is a certain kind of plan. And like plans in general, guidelines in particular are subject to a process of refinement. The process begins with the guideline description as it emanates from a body, such as the National Institutes of Health, and it concludes with the specification of specific sequences of work-orders prepared on site from one day to the next and relating to actual procedures to be performed. 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, health care personnel, and health care organizations (Wald *et al*., 2001).

Workflow systems typically employ a tripartite categorization of: *cases*, *work-items* and *resources*.[1] 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. Each case is composed of various entities, both independent – such as human beings, buildings, equipment – and dependent – such as roles and functions, as well as the processes in which these dependent and independent entities are involved. Tasks correspond to the procedures in the terminology introduced above. In the workflow context, however, tasks are not restricted to clinical procedures and not even to procedures carried out by humans, but comprehend all the kinds of activities that can be carried out by both human and non-human agents.

To understand these different classes in relation to the ontology framework outlined above, we note first of all that while *plans* and *guidelines*, as already noted, belong to the level of what is abstract, *cases*, *work-items*, and *resources* 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 activities (work-items performed by agents) and associated

[1] http://www.wfmc.org/standards/docs/TC-1011_term_glossary_v3.pdf

processes. The transition to plan formulation through to execution is captured also by the HL7 Reference Information Model v3, a developing messaging standard, which employs for this purpose various 'moods' associated with successive acts, including: Definition, Intent, Request, Event, and Goal. Here the event mood corresponds to what we are calling the execution of the act, while other moods deal with formulation, planning and assessment.[2]

The dimension of *functional role* reflect the fact that not every agent within a health care organization can carry out every task. Agents have different functional roles 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.*

5. THE ROLE OF TEAMWORK

Existing applications standardly assign work-items not to teams but to specific workers. An application for work-item assignment may in some way recognize that teams exist, but it is then 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 management 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 individuals with *collective* functional roles. They define the class *team_type*, which is instantiated by such collections of individuals. 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. Teams have collective functions which are carried out by its members working cooperatively. Team members have individual functions which overlap with one another and together form the collective functions of the team.

In what follows we describe a general methodology for modifying systems for guideline-based management of health care organizations in such a way as to allow assignment of work-items not only to individuals but also to teams. The implementation is stratified, which means that it defines the successively more refined levels at which clinical guidelines can be described in the course of implementation within a given organization.

6. PARTITIONS, VAGUENESS AND APPROXIMATION

When plans are applied in specific real-world contexs then they become refined in the course of execution. 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 the agents and teams of agents) involved in the execution of the plan in corresponding procedures.

[2] http://www.hl7.org/library/data-model/RIM/C30202/rim.htm

A relation of *refinement* can thus be defined among different plans executed on similar or overlapping sets of entities. We now show how a formal theory of this relation can be applied in the domain of workflow systems, a typical area of guideline exploitation.

Consider a map of the 50 United States, or a travel itinerary showing the several steps of a journey from JFK to LAX. We can think of each as a window upon a certain portion of spatial and temporal reality, respectively. In both cases we have a certain entity – a nation, a journey – which is divided into discrete units or portions rather as if a window is divided by a grille. If we show in addition the counties within the 50 States then the resultant map is a *refinement* of the map with which we started, which means that it manifests a more complex, hierarchical structure within which the structure of the original map is properly embedded. Each of the mentioned partitions projects onto reality in its own way, whereby each projection relation is marked always by a certain level of granularity. In each partition, too, there are factors – the soil types and vegetation existing in different areas, the precise route taken by the plane from JFK to ORD or from ORD to LAX – which fall beneath this level of granularity and thus beneath the relevant threshold of significance.

Similarly, now, in the medical domain. When prescribing or carrying out orders to administer medicines we will in effect be employings partitions of a granularity determined by dosage levels; the precise chemical structure of the prescribed pharmacologic substance will fall beneath the threshold of salience. Similarly the precise steps which need to be taken in administering the substance will also characteristically fall below the threshold of salience of a given clinical guideline or work-order, and they will thus be traced over in the granular partitions which correspond thereto. 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.) Our idea, now, is that in order to highlight different aspects of a health care organization and of its workflow processes different partitions at different levels of refinement will be needed, together with a framework within which these different partitions can be manipulated simultaneously.

To this end we call in aid the Theory of Granular Partitions (TGP) put forward in (Bittner & Smith 2003, 2003b, 2003c) and based upon foundational theories – of mereology, topology, and so forth – which are also part of the ONIONS/DOLCE approach to formal ontology. A partition, from the perspective of TGP, consists of a network of units and subunits, the latter being nested within the former. The units, in turn, are projected onto entities in reality or onto the units of other partitions. Thus, the theory of granular partitions has two parts:

A. a theory which relates units to other including units and ultimately to the partition structures in which they are included, and

B. a theory which deals with the way units and subunits of partitions are projected onto other entities for example in such a way that the subunit on the side of the projecting partition is reflected in a part-whole relation on the side of entities in reality.

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 a given health care organization which are needed for a complete ontology of its team-based workflow. When a particular human resource, for example nurse A, is entitled to carry out a particular function F, then the unit labelled nurse A in the corresponding partition of human resources is projected onto the function labelled F in the associated partition of responsibilities. When we assess how A exercises this function, then we have a new partition, where the unit labelled nurse A is projected onto the processes she actually does perform – these processes, too, being apprehended by the relevant partition always at some appropriate level of granularity.

7. REFERENCE PARTITIONS AND THE TREATMENT OF FLEXIBILITY

In the presence of team-based workflow organization 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 such flexibility by conceiving the relation between the abstract specifications contained in guidelines and the functions exercised by actual health care 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 may need to be left open in a team-sensitive framework for task-assignment. Alternatively it may 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 United States 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 an 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 unit corresponding to a given weather phenomenon in a US metereological partition can be projected onto a map of the fifty US States. (For the formal treatment of these matters in TGP see (Bittner & Smith 2003c); for the relation to the theory of Stratified Rough Sets see Bittner and Stell (2003).)

The same idea can be used, now, 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 health care 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 health care 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 required to perform the function then A's assignment projects onto the unit 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*).

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. It can also be extended to deal with those cases where the functions represented in a guideline-based workflow model are, for whatever reason, not instantiated fully accurately in actual behavior.

8. GUIDELINE-GENERATED PARTITIONS

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

1. *physical and organizational structures*,
2. *human resources*,
3. *tasks that can be assigned to the human agents within the organization,*
4. *tasks recommended by the guideline itself.*

Such partitions are organized in modular fashion. Thus the physical structure of the hospital 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 health care *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 a 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 account of partitions as structures of corresponding units and subunits. 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.

9. IMPLEMENTATION FORMALISM

Each granular partition has a root, which is, as its were, its all-inclusive unit, that of which all other units in the partition are subunits. The respective roots of the trees corresponding to the four partitions mentioned above are:

1. *hospital* (*operational units*),
2. *human agents*,
3. *tasks that can be committed to 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 (let us say) of *physician*, *resident A*, *resident B*, *nursing staff A*, *nursing staff B*, *nursing student A* and *nursing student B*. We specify four stages of conformance to the workflow in light of recommendations contained in the relevant guidelines.

G1	Specification of clinical tasks without subtaks
G2	Specification of clinical tasks with subtasks
G3	Specification of the hospital locations
G4	Specification of the health care team members, their roles and coordination

Table1. Various levels of comformance to the workflow involving recommendations from clinical guidelines

G1**:** At the coarsest level of granularity we do not consider the interior structure of a given team at all. Rather, the system records merely the fact that this 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 the ***G1*** 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 to the decision of health-care practitioners and teams at the point of care. 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 only in a coarse-grained fashion in the partitions mentioned under ***G1*** 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, functions are more precisely specified via a more refined partition. It then becomes explicit for example 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 operational unit of the hospital, the framework assigning not merely specific functions to *internal medicine team A* in relation to a given case, 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*, *surgery*, *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 not only such allocations of functions to individuals but also the cooperation between human beings in the performance of the single functions involved. Such participation may be mentioned explicitly in the guideline implementation or it may be left implicit – as something that is taken for granted by those working 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.

10. CONCLUSION

We have provided the outlines of a workflow management system which allows the representation of an ordered sequence of successively more refined modular structures (granular partitions) to support the team-based assignment of the functions specified in guidelines and the tracking of the performance of these functions by human beings in a given health care organization. This formalization is intended to be used as a basis for supplementing guideline management applications designed to help ensure that the day-to-day operations of a health care organization are in conformity with the recommendations contained in clinical guidelines. Some granular partitions consititute 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 – including the flexibility to deviate from the guideline recommendations – is indispensable if we are to build workflow systems which will be accepted by the medical community.

A further advantage of our framework is that it can enable 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 and in the context of hypertension management). This allows the true clinical situation to be addressed in a more adequate way than in existing 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 ways in which they are related together within complex, dynamic wholes. It provides also a framework for merging different task ontologies deriving from different clinical guidelines, which can be useful for example in the 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 carried out manually (though our experience tells us that domain experts who are called upon to perform such tasks very quickly become proficient at providing information in the needed 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, 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. Toward a standard for guideline representation: an ontological approach. Proc AMIA, 1999.

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.

Stefanelli M. Knowledge Management in Health Care Organizations, IMIA Yearbook of Medical Informatics 2004, 2004, 144-155. (A revised version of this review paper will appear on Methods of Information in Medicine.)

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 is supported under the auspices of the Wolfgang Paul Program of the Humboldt Foundation and also of the EU Network of Excellence in Semantic Data mining and the project “Forms of Life” sponsored by the Volkswagen Foundation.