

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

Correspondence address: Email – akumar@ifomis.uni-saarland.de

Postal – Anand Kumar (University of Saarland), Hansjakobstrasse 119 b, Munich 81825,
Germany.

Phone – +49-1728984640, +49-89-44458837

Abstract

Objective: Clinical guidelines are special types of plans realized by collective agents. We provide an ontological theory of such plans that is designed to support the construction of a framework in which guideline-based information systems can be employed in the management of workflow in health care organizations.

Method: The framework we propose allows us to represent in formal terms how clinical guidelines are realized through the actions of

individuals organized into teams. We provide various levels of implementation representing different levels of conformity on the part of health care organizations.

Result: Implementations built in conformity with our framework are 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. They do justice to the fact (1) that responsibilities within a health care organization are widely shared, and (2) that health care professionals may on different occasions be non-compliant with guidelines for a variety of well justified reasons.

Conclusion: The advantage of the framework lies in its built-in flexibility, its sensitivity to clinical context, and its ability to use inference tools based on a robust ontology. One disadvantage lies in the complication of its implementation.

Keywords : Clinical Workflow, Ontology, Clinical Guidelines, Reference Partitions, Guideline Conformance

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.

In an optimal scenario, guidelines would be integrated with the information systems operational in a given hospital in such a way. They would be sensitive to the situation of every patient in real time and able to cope, for example, where a range of different processes need to be executed in parallel on different patients in such a way as to prevent unnecessary duplication and warn of possible errors and omissions. They would be able to optimize the allocation of resources and generate work-schedules for specific individuals and groups in the organization 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

In what follows we present an ontological framework for the construction of computerized guideline support systems. Well-constructed ontologies serve not only to facilitate the re-use of already formulated knowledge in new settings, but serve also as foundation for standardization efforts, since they make explicit the tacit assumptions

underlying existing medical task flows. What is required is a reliable formal framework for the unambiguous representation of the different types of objects and processes represented in guidelines of a sort that can support computational implementations.

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).

Here each oval represents an ontology, i.e. a module that embodies the formal definitions of the corresponding types of objects and processes and of the relations between them. Arrows denote inclusion relations between ontologies, a more specific ontology using the definitions 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 types 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 for *medical procedures* comprehends also definitions of the main types of medical procedures, including for example *screening procedure*, *surgical procedure*, *diagnostic procedure*, and *laboratory procedure*.

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. This integration is currently maintained with reference to the DOLCE-Lite-Plus ontology library (Masolo *et al.*, 2004). We believe that the development of such foundational ontologies will prove to be essential for the coherent integration of heterogeneous guideline models within a single platform.

Most important for us here is the fact that ON9.3 includes ontologies of agency and plans, integrated now to the DOLCE foundational ontology and its extensions: the *ontology of descriptions and situations* (Masolo *et al.* 2004), and the so-called *ontology of plans* (Gangemi *et al.*, 2004). ON9.3 relations such as: *bearer*, *target*, and *diagnostic action* are defined within the framework of the ontology of *agency*, which concerns the entities which play particular *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* (an object pertaining to the social world, which represents the conceptualization of one or more agencies), 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, or other information objects, e.g. in the physical realization of a *diagram*.

In Figure 1, we present a UML diagram which sketches the main relations between entities in the ontology of plans (Gangemi *et al.*, 2004).

The task structure of 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 the setting of a *plan execution*, in contrast, involves only

one single temporal path of actions. A plan execution is a *situation* in the *descriptions and situations* ontology. The *situation* class also subsumes workflow “cases” (see below).

Each successive step within a task structure may be executed through a range of different actions, provided only that the constraints of the plan are satisfied. A multiplicity of different kinds of agents may be involved in the execution of a single plan. A plan can also contain *parameters* that express the required attributes or qualities that agents, resources, and actions should have in order to be compliant to the plan, at a certain granularity.

Tasks, roles, and parameters are types of *concepts defined by a plan*. Agencies, resources, actions, and attributes (*regions*) in a plan execution are *classified by* those concepts.

The *refinement* relation holds for any *social object* (plans, tasks, roles, parameters, and plan information are all taken as social objects here), and allows to talk e.g. of plans or task structures at different levels of granularity (see below).

[Figure 1]

4. THE ROLE OF CONTEXT

Stefanelli (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 a step 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 the 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.

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* here correspond to the procedures (conventional sequences of actions) in the terminology introduced above.

While *plans* and *guidelines* are social entities (representing the conceptualization of given agencies), *cases*, *work-items*, and *resources* are concrete real-world entities. When a plan is *applied* in a given health care situation (case), then the tasks are transformed

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

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 (equivalent to a plan in our ontology) is *executed*, this gives rise to a temporal sequence of activities (work-items performed by agents) and associated processes.

Different *agents* have different *roles* and thus execute different *functions (tasks)* within different contexts. Not every agent within a health care organization is authorized to carry out every task.

5. THE ROLE OF TEAMWORK

Quaglini and co-workers have formulated care flow models designed to do justice to the fact that a framework for guideline implementation should support 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 (Quaglini *et al.*, 2001).

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, teamwork – 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.

An ontology of collectives is proposed in (Bottazzi *et al.* 2004) which extends the ontologies of descriptions, situations, and plans, by providing new classes and relations that allow to distinguish collective roles and member roles, in the context of unifying plans. A team type would be an *intentional collective* in our ontology.

In what follows we concentrate on 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. For this reason, instead of describing the technical details of task structures and role hierarchies, we tackle the *refinement* relation by adopting a theory of granular partitions.

6. PARTITIONS, VAGUENESS AND APPROXIMATION

When plans are applied in a specific real-world context, 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 those involved.

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.

Compare the way in which maps can be of larger or smaller scales and can contain greater or lesser detail. Maps are windows on a certain domain of reality. They divide this domain into units or portions as a window is divided by a grille, which may be more or less *refined*. It will thus make visible certain aspects of reality but leave others invisible because they fall 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 employing partitions at the granularity of dosage levels; the precise chemical structure of the prescribed pharmacological substance will fall beneath the threshold of salience, as will the precise steps to be taken in administering the substance.

Plans in general are in every case granular in just this sense, and our idea, now, is that in order to comprehend the workflow processes of a health care organization different partitions at different levels of refinement will be needed.

To this end we call in aid the Theory of Granular Partitions (TGP) put forward in (Bittner & Smith 2003) and based upon foundational theories – of mereology, topology, and so forth – which are also part of the ON9.3 ontology library, and the DOLCE extensions mentioned above. Hence, it is straightforward to create a new extension that axiomatizes the *refines* relation according to TGP.

A partition, from the perspective of TGP, consists of a network of units and subunits, the latter being nested within the former.

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 crosscutting 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 labeled ‘nurse A’ in the corresponding partition of human resources is projected onto the function labeled ‘F’ in the associated partition of responsibilities. When we assess how A exercises this function, then we have a new partition, where the unit labeled ‘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 teamwork 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 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. We can distinguish three ways in which an indeterminate item such as a weather phenomenon can be projected onto a determinate reference partition: with full overlap (*fo*), with partial overlap (*po*), or with no overlap at all (*no*). (See Bittner and Smith (2003), Bittner and Stell (2003).)

The same idea can now be used to give 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 functions mentioned in a given guideline text are determinate, and so the latter can be used to determine reference partitions that enable us to assign functions in such a way as to do justice to the flexibility of team-based organization.

Consider the assignment of responsibilities to 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*, *allowed* or *excluded* from performing that function. 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 then allows us to calculate various combinations of assignments in a way which allows us to detect errors in granting responsibilities or in time allocations to specific individuals. Thus it can deal for example with cases 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 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 by providing an automatic means of

recovering from context-determined deviations from a given plan in the course of its realization

8. GUIDELINE-GENERATED PARTITIONS

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

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 divides the tasks of hypertension management into:

blood pressure measurement, classification of blood pressure, cardiovascular disease risk determination, etc.

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.

Here each more detailed partition is in TGP terms a *refinement* of the corresponding less detailed partition.

9. IMPLEMENTATION FORMALISM

We can now create a team-enabled workflow model for the NIH Hypertension Guideline. 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 subtasks
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

Table 1. Levels of conformance to guideline recommendations

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 organizational 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, 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. 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, etc.

Functions are now specified via a more refined partition. It has been made explicit, for example, that *blood pressure measurement* overlaps only partially with *cardiovascular disease risk determination*. 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*. *Blood pressure measurement* can appear as a subfunction of different complex functions such as *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.

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 human resources present in the health care organization will also have a modular structure, being divided into submodules under the headings: *physicians*, *nurses*, *technicians*, *laboratory attendants*, *nurse students* and so on. When we create workflow representations which do justice to the existence of *teams* of human resources this does not mean that individualized role assignments do not exist; rather they are 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 or no overlap along the lines set forth in our discussion 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*.

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, he does not directly monitor this performance and such direct monitoring is also not required by 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, flexible, modular structures 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. 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 that 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). 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 in the management of patients with multiple clinical disorders.

We are in the process of defining conformance criteria for such a system. Good candidates for such criteria are: ratio of tasks specified by a health care organization in its workflow and those present within clinical guidelines; the ability to deal both with sequentially and simultaneously occurring tasks and subtasks; the ability to deal not only with task modules carried out by teams but also to do the way in which team activities may be affected by the communication between members of a team.

One problematic feature of the framework should already be mentioned here. It is clearly highly complex, and the specification of the relevant 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).

Acknowledgment: Work on this paper is supported under the auspices of the Wolfgang Paul Program of the Alexander von Humboldt Foundation, the Semantic Mining EU Network of Excellence, The Metokis FP6 IST Project, and the project “Forms of Life” sponsored by the Volkswagen Foundation.

REFERENCES

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

Bottazzi E, Catenacci C, Gangemi A, Lehmann J, Collective Intentionality: An Ontological Perspective, C. Castelfranchi, L. Tummolini (eds.), International Conference on Collective Intentionality, to appear, 2004, also as LOA report, <http://www.loa-cnr.it>.

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

Gangemi A, Borgo S, Catenacci C, Lehmann J. Task Taxonomies for Knowledge Content, METOKIS Deliverable D07, 2004, available at <http://metokis.salzburgresearch.at>, or <http://www.loa-cnr.it>.

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.

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

Masolo C, Borgo S, Gangemi A, Guarino N, Oltramari A. WonderWeb Deliverable D18: The WonderWeb Library of Foundational Ontologies, <http://wonderweb.semanticweb.org> (2003)

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.

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

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 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.

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

