Combining guidelines and problem-solving methods for modelling therapy planning in medicine *

M. Taboada[1], D. Martínez[2], R. Marín[3], F. Palacios[4]

[1]Dpto. de Electrónica e Computación. Edificio Monte da Condesa.
Campus Sur. Universidade de Santiago de Compostela.
15782 Santiago de Compostela. Spain.

[2]Dpto. de Física Aplicada.
Universidade de Santiago de Compostela.

[3]Dpto. de Informática, Inteligencia Artificial y Electrónica.
Universidad de Murcia.

[4]Hospital General de Elche.

Author to contact: María Taboada. E-mail: chus@dec.usc.es
Telephone number: 34-981-563100
Fax number: 34-981-528012

Abstract

From the beginning of the 90's, clinical guidelines are increasingly being applied in several areas of medicine and many guideline representation languages have been proposed. These languages provide a very rich set of primitives for specifying unambiguous plans. Nevertheless, there are medical domains in which no well-established standard treatment plans exist. For these types of domains, we can model the therapy planning task by selecting an appropriate PSM and representing the domain ontology by a guideline representation language. In order to facilitate to glue PSM-based specification and guideline-based representation together, we are using the framework PROTÉGÉ and the EON system.

*This work has been funded by the Secretaría Xeral de Investigación e Desenvolvemento da Xunta de Galicia, through the research project PGIDT01-PXI20608PR, and partially by the Spanish MCyT under the project TIC2001-4936-E

1 Introduction

From the beginning of the 90's, clinical guidelines are increasingly being applied in several areas of medicine, including therapy planning. The most direct benefits of their application are to improve patient care, to reduce practice variability and to decrease patient care costs [9]. These guidelines are in narrative form, so their implementation in a patient-care context by incorporating patient-specific data could substantially reduce costs through undesirable practice variation [13]. One alternative to "static presentation" of guideline knowledge is to generate recommendations by monitoring or checking the consistency of the medical actions. Currently, many research groups are developing sharable guideline representation languages designed for modelling guidelines in terms of primitives representing actions, decisions, eligibility criteria and plans. Some representative approaches are GLIF [14], EON [11] or PROforma [7]. However, we argue that there are medical domains in which no well-established standard treatment protocols exist, and the physician has to decide on the therapy that is to be applied to each patient, in function of a set of therapeutic objectives to be fulfilled. For this reason, we propose the modelling of this type of task by combining a guideline representation and some standard algorithm (problem-solving method, PSM) designed for automating this task, such as the proposed in [19]. The representation of the unambiguously defined parts of the protocol can be efficiently made by using a guideline representation language, whereas the rest of the protocol can be modelled by problem-solving methods (PSMs). The current development of a PSM-based decision support system can be view as consisting of two classes of reusable components [10]: domain-independent PSMs (or standard algorithms for automating stereotypical tasks) and domain ontologies (representation of domain concepts and relationships among concepts). Combining the guideline-based representation and the PSM-based development requires a component-based representation framework that facilitates the integration among knowledge structures representing guidelines and PSMs, such as the EON system [11]. In this paper, we are presenting a medical example which integrates guidelines and PSMs.

2 Modelling treatment administration in medicine

Treatment administration has been often modelled as a 'planning protocol-directed therapy' task. Oncocin [21] and, later, the T-Helper system [20] were some of the first therapy-advice systems. Oncocin was designed following a standard algorithm for automating the planning task known as 'skeletal-plan refinement' [8]. The T-Helper system implemented a protocol for AIDS treatment using the PROTÉGÉ framework [16]. Its architecture, the EON framework [11], was independent of the application domain. In this way, the AIDS knowledge base could

be substituted with another knowledge base. Currently, other many guideline representation languages have been proposed, such as GLIF [14] or PROforma [7]. These languages provide an extensive set of modelling primitives oriented to represent actions, decisions and plans.

Recently, attention has been paid to formal model guidelines. For example, the library of ontologies ON9.2 [15] includes ontologies related to guidelines. In this library, a guideline is modelled as a kind of plan and it is represented by a flowchart. So, the structural part of the guideline (the flowchart) is explicitly separated from the semantic part (the plan). In addition, we can consider a third level related to a guideline, the level of the underlying procedure. In therapy planning, the 'skeletal-plan refinement' algorithm is a very used reasoning strategy. This method consists of selecting a general plan of predefined treatment (called a 'skeletal plan'). Each plan specifies standard treatment protocols, and their modification in function of the patient's state. Once a plan has been selected, it is refined cyclically into ever-more detailed plans, until a complete plan that is applicable to the patient is obtained. The plan obtained determines all the possible sequences of treatments that may be applied to the patient undergoing a protocol. The proposed guideline representation languages provide a very rich set of primitives for specifying unambiguous plans.

Nevertheless, the 'episodic skeletal plan refinement' algorithm cannot be applied in medical domains in which no well-established standard treatment protocols exist. This is, for example, the case of Intensive Coronary Care Units (ICCU). For these types of domains, we have proposed to model the task of treatment administration as a design task [19]. The generic models for the design task proposed in the CommonKADS library [18] correspond to the types of problems referred to by Chandresekaran and Johnson [2] as 'class 1, class 2 and class 3' design problems. In CommonKADS, these types of problems are referred to as 'creative design, innovative design and routine design', respectively. In class 3, or routine design, the manner of breaking down the problem, as well as the design plans to be applied in each stage, are known beforehand. Therefore, this problem is analogous to 'episodic skeletal-plan refinement'. In class 2, or innovative design, the components are known beforehand, but not the design plans, and in class 1 design, neither the components nor the plans are known. Our task is adapted to class 2 design, since although the physician selects a treatment from amongst the drugs of which he has information, he does not have a detailed plan to determine which specific drug should be applied to the patient at each moment in time. For this reason, we can adapt the general class of algorithms for design that Chandrasekaran [1] calls 'Propose-Critique-Modify' (PCM).

3 Combining the guideline-based representation and the PSM-based development

From the 80's decade, when Allen Newell introduced the 'knowledge-level hypothesis' [12] and Clancey's and Chandrasekaran's research groups [2, 3] demonstrated that recurring problem-solving strategies were the core of many intelligent systems, many methodologies were developed for designing intelligent systems. Some of these are CommonKADS [18], Protégé [16] o UPML [6]. Following these methodologies, the design of an intelligent system can be efficiently automated using two classes of reusable building blocks:

- Domain ontologies, which describe the concepts and relationships among those concepts for a specific domain. For example, nowadays in the medical domain some ontologies are available on-line, such as the library of ontologies ON9.2 [15], which includes both general and specific ontologies, the UMLS semantic network ('www.nlm.nih.gov') , which provides a medical ontology developed for assisting information retrieval or the ontology provided by the EON System [11], which describes the general structure of clinical protocols.

- Problem-solving methods (PSMs), which are domain-independent standard algorithms oriented to automate stereotypical tasks. One of the most known PSMs is the 'heuristic classifier' proposed by Clancey [3], which breaks down a classification task in three subtasks: abstraction of data in general features, heuristic matching of general features to a set of possible solutions and refinement of general solutions into more specific solutions. Each PSM clarifies what domain knowledge is needed and what is the purpose of this knowledge to solve a particular task. In the example, abstraction knowledge, relationships between general features and solutions and solution hierarchies are the knowledge required.

Currently, the development of an intelligent system can be viewed as the assembly of domain ontologies and PSMs [10]. The main steps in the development of an intelligent system are three [5]:

- Task analysis, which identifies the problem and obtain the I/O relationships and the available knowledge.

- PSM selection, which identifies a suitable problem-solving method.

- PSM configuration, which links the PSMs to a domain ontology.

Following this perspective, an intelligent system is viewed at a higher level of abstraction, as a consequence of separating domain concepts from the use of these concepts. As a result of this, the direct reuse of previously tested solutions

is possible and the development and maintenance of each new system is easier [10].

On the contrary, a guideline is difficult to reuse as domain and procedural knowledge are often intertwined. In addition, in some medical domains no well-established treatment protocols exist, so the therapy planning task do not easily adapt to the format used in a guideline. In spite of that, we think that the use of some guideline representation language has many advantages as they provide a very rich set of modelling primitives.

The alternative proposed here is to model the therapy planning task by selecting an appropriate PSM and representing the domain ontology by taking into account, as much as possible, a guideline representation language. In order to facilitate to glue PSM-based specification and guideline-based representation together, we are using the PROTÉGÉ framework [16] and the EON system [11]. The first is a tool oriented, among other things, to develop domain ontologies and to facilitate the reuse of knowledge. The last feature is reached by allowing to incorporate information directly from UMLS and to merge previously defined ontologies into your domain ontology. On the other hand, the EON guideline modelling language allows expressing different types of subtasks included in a plan [22], such as setting goals, choosing among different alternatives, sequencing actions and interpreting data. In EON, guideline modelling uses a set of interacting models: 1) a core model that defines the set of concepts and relationships among these concepts used in guidelines, and 2) a set of models containing modelling primitives in order to represent different types of knowledge in guidelines. For a specific application, all these models must be specialized and assembled. From this point of view, the development of electronic guidelines is viewed as the assembly of several models in a similar way to the current design of intelligent systems.

4 An example

This section describes an example showing the modelling of the treatment administration task in the clinical domain of Intensive Coronary Care Units (ICCU). We view treatment administration as the task of planning a therapy and monitoring a patient following that therapy. The tasks of therapy planning and patient monitoring have been identified as being two generic tasks in medicine [17, 23]. The objective of the former is to choose the most adequate treatment for the state of the patient, given a set of possible treatments. In our model, the input data for this task are the results of the diagnostic process and the relevant characteristics of the state of the patient. These data are abstracted in a set of therapeutic objectives which orientate the therapeutic choice. Conversely, the monitoring of patients consists of observing and controlling the state of the patient, taking into account the disorders diagnosed and the therapy applied. During the monitoring

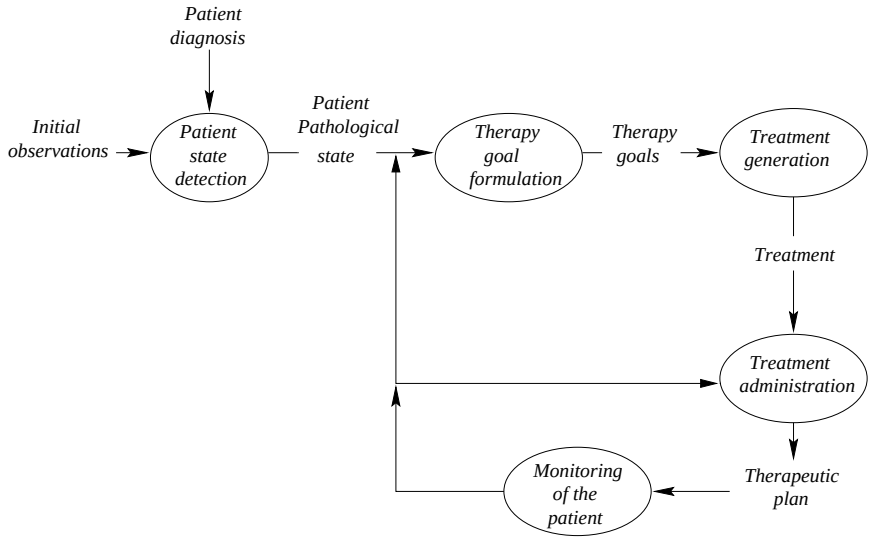

Figure 1: Data flow in the treatment administration task

of the patient, deviations from the expected state of the patient can be detected, which force a readjustment of the treatment or therapeutic objectives. These readjustments usually give rise to new therapy planning.

In summary, these two medical tasks consist of a cycle in which the state of the patient is continually evaluated, and in function of this, the therapy to be applied is revised. In the ICCU domain, the manner of carrying out these two tasks gives rise to the cyclical execution of the following subtasks (Fig. 1):

1. *Patient state detection*, which seeks the set of initial observations the values of which deviate from those expected (called abnormality observations). Both of the sets of abnormality and normality observations are abstracted in the *patient pathological state*. That is, the pathological state of a patient is described by a set of high level observations, such as 'low contractility, high preload and compensated afterload'.

2. *Therapy goal formulation*, based on the patient's pathological state. These goals are expressed by the physician as 'qualitative changes' of high level observations. For example, for the pathological state of patient previously described, the set of therapeutic goals will be 'to improve contractility, to reduce preload and to maintain afterload'.

3. *Treatment generation*, which selects a set of drugs, the joint administration of which permits achieving the formulated therapy goals.

4. *Treatment administration*, generating the dosage and temporal pattern for administering each drug of the treatment, taking into account the patient's pathological state.

5. *Monitoring of the patient*, detecting possible complications that may arise due to the administration of the treatment. The detection of complications can give rise to the formulation of new therapeutic goals or to the administration of a new treatment.

Taking into account the main substasks, we can model the treatment administration task by adapting the generic class of problem resolution methods for design task, which Chandrasekaran [1] labels Propose-Critique-Modify (PCM). In this type of PSM, the design task is broken down into four subtasks: proposal of a solution, verification of the proposed solution, critique of the solution, and modification of the proposal. Patient state detection, therapy goal formulation, generation of a treatment and treatment administration are proposal-style subtasks. The *monitoring of patients* is carried out in three subtasks: *evaluation of the patient's response to the treatment*, *expected complication detection* and *modification of the treatment*. The initial two are verify-type subtasks, and the third one is modify type. The justification for the absence of a critique type subtask is as follows. A critique type subtask would require complete knowledge of how each drug acts on the patient. Nevertheless, at present, for the majority of these drugs, these models have not yet been developed, and those that do exist are model proposals that have not yet been verified.

4.1 Patient state detection

This task corresponds to the generic modelling of system *monitoring* . This model compares the real observations to be controlled with the optimum values which it is desirable to obtain. The discrepancies thus obtained are classified into qualitative discrepancies. Thus, the most appropriate manner for resolving this task is by means of a simple classification PSM. The knowledge required by this PSM can be represented in several ways on EON. In the clinical domain of the example, a kind of haemodynamic classification for a patient suffering from acute myocardial infarct is the Forrester's classification. Each grade (I, II, III and IV) of Forrester's haemodynamic classification for a patient can be specified by a Boolean combination of other criteria, and a single criteria can be specified as a numeric term criterion, a presence criterion or a temporal criterion. For example, Fig 2 shows two PROTÉGÉ forms specifying the comparison of the numeric concept 'Cardiac Index (CI)' with its lower limit of normal value and the definition of 'High Preload' in terms of the value of 'Artery wedge pressure (PWP)', respectively. Fig. 3 combines the two previously defined criteria in order to classify the diagnosis of the patient in terms of the Forrester's grade.

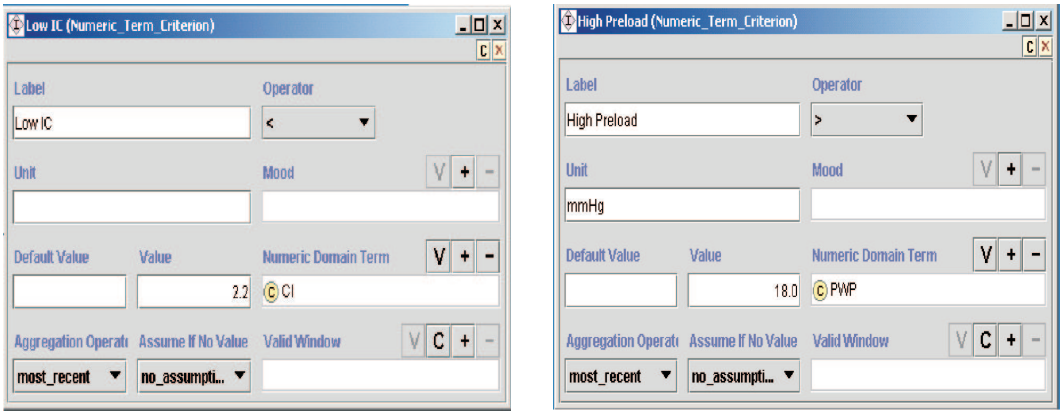

Figure 2: Specification of two numeric term-based critera using EON

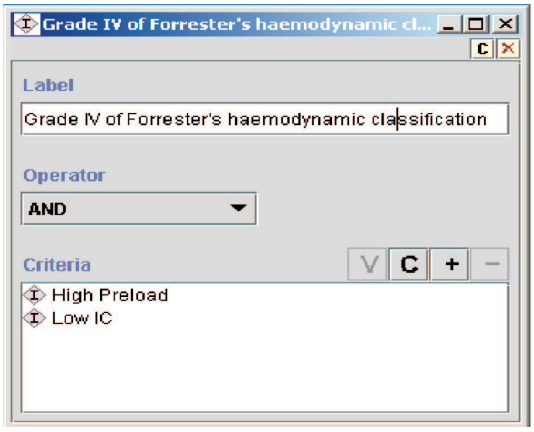

Figure 3: Specification of a boolean combination of two criteria using EON

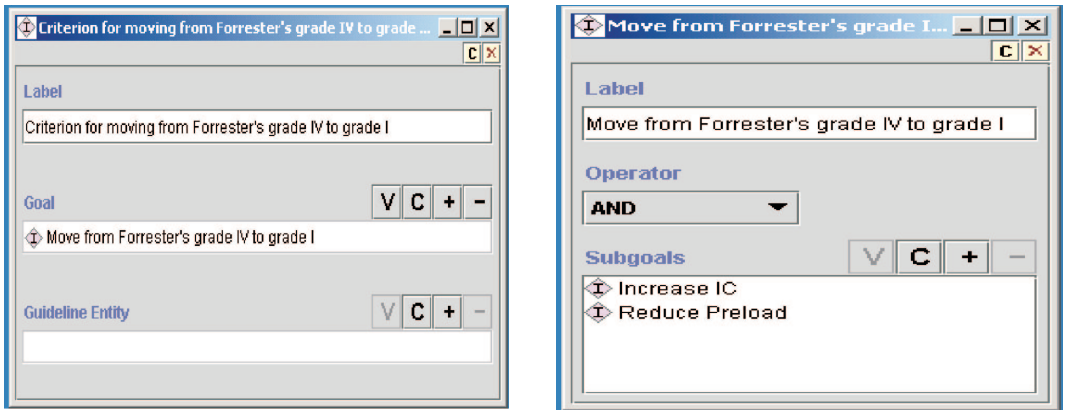

Figure 4: Specification of a goal criterion and of a composed goal using EON

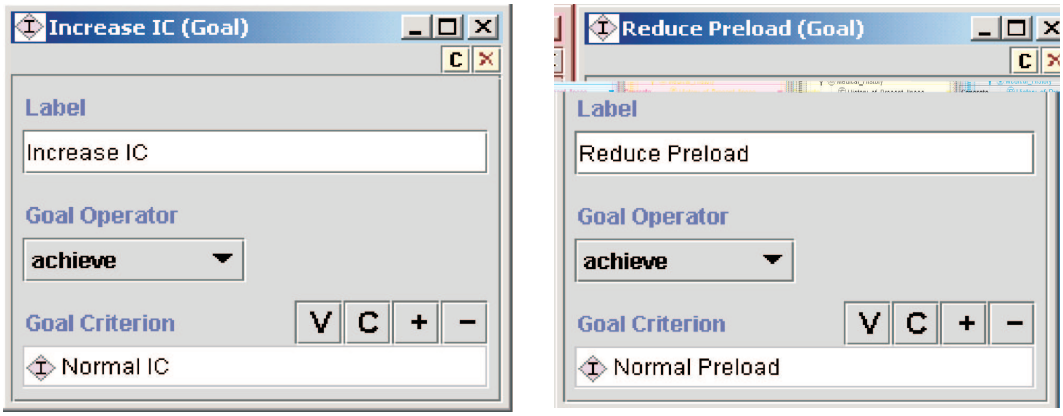

Figure 5: Specification of two single goals using EON

4.2 Therapeutic Goal Formulation

The *Therapeutic Goal Formulation* task only contains one inference, which transforms the patient's pathological state into a set of therapeutic objectives to be attained. This is again a classification PSM. The knowledge required by this PSM for this subtask can be represented by the class 'goal criterion' defined on EON. In figure 4 we can see an instantiation of a goal criterion for representing the goal of moving from one Forrester's grade to another. In this particular case, the goal is composed of two single goals: 'Increase IC' and 'Reduce Preload'. Both goals have been defined by means of the operator 'achieve', such as it is shown in Fig. 5.

4.3 Generation of a Treatment

This task consists of the abduction of a treatment that is compatible with the therapeutic objectives to be obtained. The effects that a particular treatment may bring about can be forecasted by tracing pathways in causal relationship. This method requires:

- Knowledge of the domain which relates each family of drugs with the primary effects that they produce and with the inclusion/exclusion criteria that have to be borne in mind.
- The availability of the current state of the patient at significant time points, in order to be contrasted with the desired effects.

4.4 Treatment administration, evaluation and detection of expected complications

Once a treatment has been selected by the user, the next stage is that of determining, for each drug of the treatment, the therapeutic phase to be applied (temporal guidelines, dosage, etc.). The knowledge required by this task is efficiently represented by a management diagram on EON.

4.5 Modification of treatment

This task consists of modifying the treatment, either by associating or eliminating drugs, or by triggering new therapeutic objectives. This subtasks is similar to the therapeutic goal formulation subtask.

4.6 Discussion

Currently, many guideline representation languages are available. These languages provides a very rich set of modelling primitives, such as actions, decisions and plans. Nevertheless, there are medical domains in which no well-established standard treatment protocols exist. This is, for example, the case of Intensive Coronary Care Units (ICCU). For these types of domains, in [19] we have proposed to model the task of treatment administration as a design task adapting the propose-critique-modify PSM [2]. Due to the complexity and richness of the medical knowledge, we are not building the domain ontology from the scratch, rather we are representing our domain by using the guideline representation language developed in the EON project [11]. In order to facilitate to glue PSM-based specification and guideline-based representation together, we are using the PROTÉGÉ framework [19] and the EON system [4]. In this way, the modelling of the treatment administration task can be viewed as the assembly of the domain model and the propose-critique-modify PSM. Each particular domain model can

be developed by instantiating the general ontology defined by the EON system. But in order to extend the EON ontology, we should reuse knowledge, as much as possible, for achieving a sharable system. We have available some tools oriented to merge or incorporated ontologies and terminology servers, but the reuse of knowledge is still difficult. We tend to introduce new concepts and relationships when we do not quickly find them and this practice runs into reusing. Until now, we cannot be sure that reuse pre-existing knowledge facilitates the development of knowledge bases [4].

On the other hand, although reusing PSMs is the recommended paradigm in the knowledge engineering area, nowadays it is difficult to put it into practice as:

- There are only available on-line a few PSM libraries.
- Until now there is no agreement on the optimal way for bringing domain ontologies and PSMs together. In many practical cases, they are connected in an implementation dependent way.

References

[1] B. Chandrasekaran. Design problem solving: A task analysis. *AI Magazine*, 11(4):59–71, 1990.

[2] B. Chandrasekaran and T.R. Johnson. Generic task and task structures: History, critique and new directions. In Krivine David and Simmons, editors, *Second Generation Expert Systems*, pages 232–272. Springer-Verlag, New York, 1993.

[3] W.J. Clancey. Heuristic classification. *Artificial Intelligence*, 27:289–350, 1985.

[4] P. Cohen, V. Chaudhri, A. Pease, and R. Schrag. Does prior knowledge facilitate the development of knowledge-based systems? In *Proc. of AAAI-99*, 1999.

[5] H. Eriksson, Y. Shahar, S.W. Tu, A.R. Puerta, and M.A. Musen. Task modeling with reusable problem-solving methods. *Artificial Intelligence*, 79(2):293–326, 1996.

[6] D. Fensel, R. Benjamins, E. Motta, and B. Wielinga. UPML: A framework for knowledge system reuse. In *Proceedings of the International Joint Conference on AI (IJCAI-99)*, Stockholm, Sweden, July 31 - August 5, 1999.

[7] J. Fox, N. Johns, and A. Rahmanzadeh. Disseminating medical knowledge: the proforma approach. *Artificial Intelligence in Medicine*, 14(1-2):157–182, 1998.

[8] P.E. Friedland and Y. Iwasaki. The concept and implementation of skeletal plans. *Journal of Automated Reasoning*, 1:161–208, 1985.

[9] J.M. Grimshaw and I.T. Russel. Effects of clinical guidelines on medical practice: a systematic review of rigorous evaluation. *Lancet*, 342:1317–1322, 1993.

[10] M. Musen. Moderm architectures for intelligent systems: reusable ontologies and problem-solving methods. In *Proc. of AMIA Annual Symposium*, 1998.

[11] M. Musen, S. Tu, A. Das, and Y. Shahar. Eon: A component-based approach to automation of protocol-directly therapy. *JAMIA*, 3:267–388, 1996.

[12] A. Newell. The knowledge level. *Artificial Intelligence*, 18:87–127, 1982.

[13] J. Overhage, W. Tierney, X. Zhou, and C. McDonald. A randomized trial of corollary orders to prevent of omission. *JAMIA*, 4(5):364–375, 1997.

[14] M. Peleg, A. Boxwala, O. Ogunyemi, Q. Zeng, and S. Tu. Glif3: The evolution of a guideline representation format. In *Proc. of AMIA Annual Symposium*, 2000.

[15] D.M. Pisanelli, A. Gangemi, and G. Steve. Towards a standard for guideline representation: an ontological approach. *JAMIA*, 6:906–910, 1999.

[16] A. Puerta, J. Edgar, S.W. Tu, and M.A. Musen. A multiple-method knowledge-acquisition shell for the automatic generation of knowledge-acquisition tools. In *Proc. of the 6th Knowledge Acquisition for Knowledge-Based System Workshop*, Banff, Canada, 1991.

[17] M. Ramoni, M. Stefanelli, L. Magnani, and G. Barosi. An epistemological framework for medical knowledge-based systems. *IEEE Trans. on Systems, Man and Cybernetics*, 22(6):1361–1375, 1992.

[18] A. Schreiber, H. Akkermans, A.A. Anjewierden, R. Hoog, N.R. Shadbolt, W. Val de Velde, and B. Wielinga. *Engineering and managing knowledge. The CommonKADS methodology.* The MIT Press, 1999.

[19] M. Taboada, M. Lama, S. Barro, R. Marín, J. Mira, and F. Palacios. A problem-solving method for 'unprotocolised' therapy administration task in medicine. *Artificial Intelligence in Medicine*, 17:157–180, 1999.

[20] S.W. Tu, H. Eriksson, J.H. Gennari, Y. Shahar, and M.A. Musen. Ontology-based configuration of problem-solving methods and generation of knowledge acquisition tools: the application of PROTÉGÉ-II to protocol-based decision support. *Artificial Intelligence in Medicine*, 7:257–290, 1995.

[21] S.W. Tu, M.G. Kahn, J.C. a, M.A. andFerguson, E.H. Shortliffe, and L.M. Fagan. Episodic skeletal-plan refinement based on temporal data. *Communications of the ACM*, 32(12):1435–1455, 1989.

[22] S.W. Tu and M. Musen. A flexible approach to guideline modelling. In *Proc. of AMIA Annual Symposium*, 1999.

[23] G. van Heijst, G. Lanzola, A. Th. Schreiber, and M. Stefanelli. Foundations for a methodology for medical KBS development. *Knowledge Acquisition*, 6:395–466, 1994.