Design of a Goal Ontology for Medical Decision-Support

Davide Zaccagnini, MD[1], Aziz A. Boxwala PhD[2,3]

[1]Language and Computing, N.V., Boston, MA

[2]Decision Systems Group, Brigham and Women's Hospital, Harvard Medical School, Boston, MA

[3]Eclipsys Corporation, Boston, MA

ABSTRACT

Background: Various computable representations of clinical decision-making have been proposed in the last few years. This study uses an ontological approach for formally representing healthcare providers' intentions, as key parts of the clinical reasoning activity. We used the resulting ontologies, to simulate the provision to health care providers of clinical plans of action based on clinicians' goals.

Methods: We developed an ontological representation of medical goals, plans, clinical scenarios and other relevant concepts in medical decision-making. Ontologies were developed in the Web Ontology Language (OWL) and organized in a distributed fashion. Examples of clinical plans were taken from two established clinical guidelines. Nine test cases were created to verify the model's ability to appropriately and exhaustively provide recommendations taken from the guidelines. For all test cases, plans were successfully identified and retrieved.

Conclusion: In preliminary tests, the ontological design that we developed supported effective reasoning over formally represented medical knowledge. The study highlighted various desirable aspects of this approach: easy management of knowledge, the possibility of exploiting assorted knowledge bases containing recommendations from disparate sources. The limitations of the study reside in the simulated setting used for testing and the lack of independent evaluation of the results.

The immediate extension of this approach into real medical information systems is limited by the lack of reliable, "industrial strength" knowledge management tools.

Background

Clinical Decision Making

The Evidence-Based Medicine[1] movement has led in the last two decades to the development and organization of vast amounts of medical knowledge in structured documents, known as clinical practice guidelines. These guidelines are intended to represent evidence-based, best practices of care for various clinical situations. Several computer-based decision-support methods have been developed and evaluated for delivery of guideline recommendations at the point of care. The current challenge is to support complex clinical decision making in real clinical scenarios where 1) patients may present with multiple disease and be engaged in various medical interventions, and 2) multiple actions, carried out at various levels of care may be required.

In such complex scenarios, selection of patient management goals appropriate to the current context is important. Formal definition of goals can be used to restrict the provision of clinical recommendations to only those in line with the clinician's current aims, avoiding irrelevant or inappropriate suggestions. In case of multiple diseases, for instance, this mechanism can focus decision support on the recommendations that are appropriate for the type of care provider or for the most urgent health issue. Clinical decisions at different levels can also be integrated in the clinical process by means of formalized goals. Managing one patient's condition frequently requires the execution of different clinical tasks. The declaration of specific sub-goals would drive decision support to the single steps, molding the guideline on the clinical workflow. This can possibly increase the adoption of clinical guidelines.

Our approach, rather then representing the recommendations in an algorithmic manner as sequence of clinical steps, defines each piece of information in a modular way, leaving the selection of the recommendations to be provided to the inference engine. The reasoning engine bases its inference on the clinician's goals and the patient's clinical scenario.

Intentional Planning

Formal definition of goals is a very active area of research in robotics and autonomous agents[2,3]. Limited work has been done in the medical domain on these topics, i.e. intentional planning, using the ontological approach, formal semantics and description logic based reasoning.

We used Bratman's philosophical framework[4], subsequently formalized in the Belief Desires Intentions (BDI)[5] model, to translate the high-level concept of goal into an ontological class, along with other relevant concepts.

Bratman defines an intention, a goal, as an agent's commitment to execute one or more plans in order to achieve a state of the world. While desires are unrelated to the agent's context (i.e., to discharge a very sick patient) and do not lead to plans selection and execution, no intention is such if a viable plan cannot be identified for its achievement. The need of instantiating only practically achievable intentions is obvious in clinical medicine where, for instance, curing a terminal illness with no known cure should not be considered as a clinical goal.

Beliefs are those facts deemed true by the agent. We identified these as the available medical information and named them the clinical state.

Goals are transient entities; as plans are executed, patient's and caregiver's environments change, and new intentions arise while old ones decay.

Material and Methods

We used a multi-partitioned representation of the targeted domain based on the BDI framework discussed above. Chances of classification errors due to multiple inheritance were minimized using description logic-based automated classification.

In the current architecture, three distinct ontologies (the Goals Ontology, the Plan Library and the Test Ontology) were created. The Goals Ontology is the core representation schema. It contains the definitional properties of high-level classes (such as Goal, Plan and Clinical_State). The Plan Library contains recommendations from two clinical practice guidelines represented as instances of the Plan class. The Test Ontology contained synthetic cases used to test the overall representation. This design reflects how we envision the operations of intentional planning, that is, as the execution of a query against a knowledge base, the Plans Library in this case.

We used Protégé ontology editor[6] and the OWL plug-in[7] as development environment. The OWL plug-in supports the representation of ontologies in the Web Ontology Language (OWL)[8].

The Goal Ontology

We defined a goal as an ordered tuple of two components, < DCS, P>, where:

DCS is the desired clinical state

P is the goal's priority

DCS is the clinical state that a care provider would like to achieve for his patient. Examples of DCS are 'normal blood pressure', in a hypertensive patient or 'excised breast lump', in a case with diagnosed breast neoplasm. P has not been further defined in the current version because of its dependency on the specific implementation of the ontology.

A plan, as already discussed, is an action or a set of actions that can be taken to achieve a given state of the world. A plan can be any process described by the following set of ordered attributes:

< ICS, FCS, C, G, CC, R>, where:

ICS is patient's initial clinical state to which the plan applies.

FCS is the final clinical state achieved by the successful execution of the plan. DCS, ICS and FCS are all instantiated using individuals of the Clinical_State class mentioned above; whether being sought, currently existing, or the result of a plan execution.

C is the plan's cost, a general function that can be further defined according to specific implementation's requirements.

G is the plan's goal and takes an instance of Goal as value.

CC is the context of care. This attribute specifies the clinical environment in which the plan can be executed, the emergency room or a small clinical practice for instance.

R defines the resources needed to execute the plan. Clinicians, medical devices, and facilities employed are all examples of resources.

Note that the Plan class does not describe the details of the actual recommendation. This is intentional since the scope of our work is limited to selection of recommendations based on the clinician's goals and the patient's clinical status. Other representations of clinical guidelines[9,10,11] provide detailed representations for descriptions of plans. The goal ontology we describe can be used with those representations.

An overview of the above classes and relationships among them is shown in Figure 1 in the form of a class diagram in the UML[12] notation.

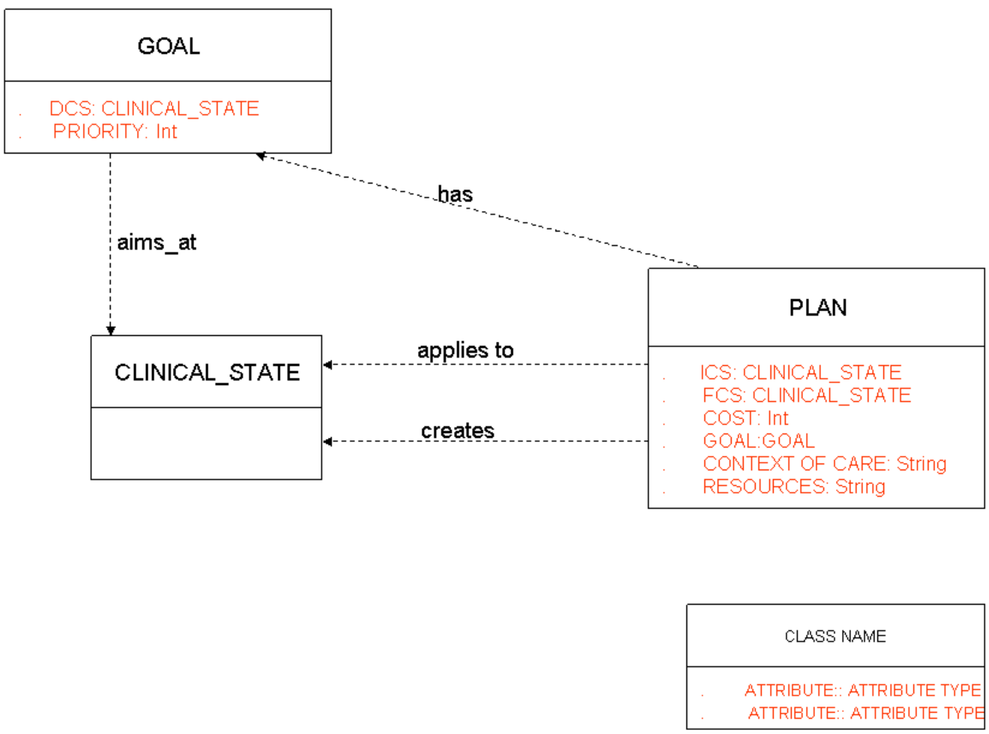

Figure 1. Goal, Plan and Clinical State classes and their relationships in the ontology.

The Plans Library

We built the plan library by modeling two clinical guidelines: diagnosis of breast disease[13] and hypertension diagnosis and treatment[14]. Both these guidelines are published by the Institute for Clinical Systems Improvement (ICSI) in Bloomington, Minnesota. The plans were constructed using the Hierarchical Task Decomposition (HTD[15]) approach. The HTD method involves subdividing complex processes down to their atomic components, defined as non-divisible action. Every action thus creates the needed conditions for the following action to start, frequently referred as plan's preconditions (Figure 2). These are represented by the ICS attribute in our design.

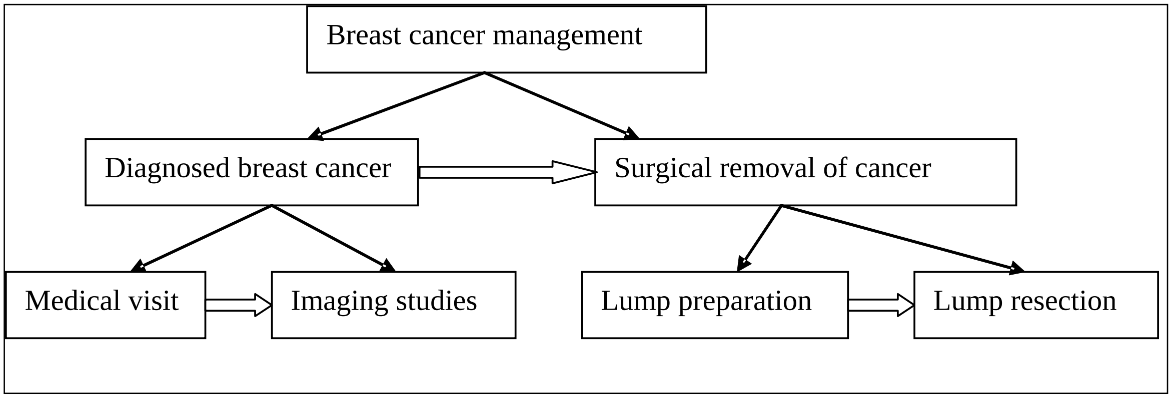

Figure 2. Hierarchical task decomposition of a complex medical plan. The thin arrows point to plans' components; the thick, hollow arrows indicate the sequence of execution of the plans'.

- **Plans for Breast Disease**
  - Plans for Breast Disease Diagnosis
    - *mammography*
    - *ultrasounds*
    - *cyst aspiration*
    - *ask about personal history of ductal hyperplasia with atypia*
    - …
  - Plans for Breast Cancer Treatment
    - *surgical resection*
    - *chemotherapy*
    - …
- **Plans for Hypertension**
  - Plans for Hypertension Diagnosis
    - *blood pressure measurement*
    - *repeated blood pressure measurement*
    - *12-lead electrocardiogram (ECG)*
    - *Urinalysis*
    - *blood glucose test*
    - …

Figure 3. Examples of hierarchically decomposed plans (atomic plans are in italics) from the Plan Library.

An atomic plan is not to be considered such in all circumstances. A radiologist may want to subdivide an atomic plan further, according to his area of expertise. We decided to limit granularity of the plans to the level of detail that would be needed by a primary care physician or a general practitioner to manage his patients.

The set of plans represented in our library do not encompass the entire set of recommendations from these two guidelines. Thirty-two recommendations were encoded for testing purposes, 22 for diagnosis of breast disease, and 10 for hypertension diagnosis and treatment.

After their sub-division, plans were formally defined using description logic statements. An example defining the screening mammogram plan is shown in Figure 4. This example and others shown in this paper are simplified. Realistic examples from the medical domain would be more complex than the examples shown here.

Plan Λ

∃hasInitialClinicalState ((∃NoPalpableMass ⊔ ∃NoDiagnosis) ⊓ ∀moreThan35years))

⊓ ∃ hasGoal DiagnoseBreastDisease

Figure 4. Formal definition of the screening mammogram plan. The statement defines a plan applicable to a patient who is older than 35 years with no palpable breast mass at the physical examination, and no established diagnosis of breast disease. The plan's goal is to diagnose breast disease.

The Test Ontology

Reasoning

The design was tested using a simulation of retrieval of plans from the Plan Library. The simulation envisioned a usage scenario involving a health care provider interested in obtaining evidence-based recommendations to manage his or her patients. We used Racer[16], a description logic-based reasoning engine, to perform inference on the ontologies.

The test ontology contained nine test cases representing patients with different presentations, modeled in the system as utility classes. The retrieval process was simulated by automatic classification of a test case under the appropriate clinical plans. As an example, the test case (Case N. 1) was instantiated by the following set of constraints:

hasInitialClinicalState((NoPalpableMass) ⊔ (NoDiagnosis))

hasAge(moreThan35years)

hasGender(female)

This case is classified, as shown in Figure 5, under the screening mammogram plan (defined in Figure 4).

Original hierarchy

- Plans
  - PlansFroBreastDiaseseDiagnosis
    - mammography
    - ultrsounds
    - cystAspiration
    - …..
- ……
- TestCase
  - Case N. 1
  - Case N. 2
  - Case N. 3
  - Case N. 4
  - …..

Inferred hierarchy

- Plans
  - PlansFroBreastDiaseseDiagnosis
    - mammography
      - Case N. 1
    - ultrsounds
    - cystAspiration
    - …..
- ……
- TestCase
  - Case N. 2
  - Case N. 3
  - Case N. 4
  - …..

Figure 5. The inference performed by reasoned and its results. Case N. 1 is correctly assigned to the screening mammogram.

The content and structure of automatically inferred hierarchies were reviewed against original guidelines and medical expertise for completeness and correctness of retrieval.

Instantiated ICS and G classes were set to be necessary and sufficient conditions in the test cases, thus defining the specific clinical scenario and clinician's intentions for that case. Other attributes were also evaluated (patient characteristics, available resources etc).

Case #	**Clinical Presentation**
1	ICS: no breast-related or other symptoms, patient older than 35 years G: screen patient for breast cancer
2	ICS: no breast-related or other symptoms, patient younger than 35 years G: screen patient
3	ICS: palpable breast mass present, no other symptoms G: establish diagnosis of breast cancer
4	ICS: possible diagnosis of breast cancer, residual mass after cyst aspiration G: establish diagnosis
5	ICS: established diagnosis of breast cancer, imaging studies completed G: treat breast cancer
6	ICS: high blood pressure values at first measurement, no other symptoms G: establish diagnosis
7	ICS: high blood pressure at repeated measurements, suspected primary hypertension G: confirm diagnosis

8	ICS: diagnosed secondary hypertension G: treat secondary hypertension
9	ICS: controlled primary hypertension, high blood glucose G: investigate high blood glucose

Table 1. Medical test cases with instantiated attributes

Plans were retrieved correctly and exhaustively in all test cases. In Case N.2, the patient's age makes her ineligible for screening or other procedures and therefore no plans are retrieved.

Cases	**Retrieved plans**
Case N.1	1) perform screening mammogram
Case N.2	2) no plans retrieved
Case N.3	1) ask about personal history of ductal hyperplasia with atypia 2) ask about personal history of ductal hyperplasia without atypia 3) ask about first degree relatives with breast disease 4) ask about second degree relatives with breast disease 5) ask about familiar history of other neoplasm 6) perform cyst aspiration
Case N.4	1) referral to specialist 2) perform ultrasound-guided cyst aspiration
Case N.5	1) referral to a specialist 2) perform surgical resection
Case N.6	1) perform repeated blood pressure measurements
Case N.7	1) test 12-lead electrocardiogram (ECG) 2) test blood glucose 3) perform urinalysis 4) test hematocrit 5) test serum electrolytes 6) test serum creatinine
Case N.8	1) refer to specialist
Case N.9	1) refer to specialist 2) refer to appropriate guideline

Table 2. Test cases and the results of plans classification

Discussion

The importance of definition of goals in medical decision-support can be illustrated by the following example: a patient with an acute cardiac insufficiency who is also obese, would instantiate in a clinical information system, a clinical scenario consisting of two conditions, obesity and acute myocardial insufficiency. However, providing dietary recommendations for this patient would be inappropriate during the immediate care of the patient because of the much higher urgency of the cardiac condition. Specifying that the current focus, the goal, is restoring the patient's normal cardiac function would avoid this undesired behavior.

The approach used in this study could also contribute to the maintainability of the heterogeneous medical knowledge bases. Multiple inheritance along with description logic-based classification, can prove to be powerful tools for managing complex domain knowledge.

Goals formalization has not often been the primary focus of research studies on computerized medical decision-support. Although other systems, primarily, but not only, ProForma[10] and Asbru[9], support specification of goals, no other study experimented their use as the central component of a decision support system. We analyzed goals with an ontological approach and grounded our representation in the Bratman's general framework for agency. The value of Goal's attribute P (priority), although not investigated in this study, is instrumental for establishing precedence among conflicting or concurring plans. In the example presented above, this attribute would be used to describe weight-reduction recommendations as still applicable, but with lower priority than those for the cardiac deficit and displayed accordingly. A recent study confirmed the strong need for a filtering methods for prioritize clinical recommendations in complex clinical cases[17].

The representation of plans (often as algorithmic flowcharts) in other systems mentioned above is certainly much more extensive than the rudimentary one that we have used. Plans representation was in fact not the first aim of our study. However, we believe that our approach to representing goals can be used to extend the other representations for medical decision-support.

In a previous study[18], we proposed a richer definition of goals encompassing five attributes rather than the current two. In the current version, the *initial* clinical state has become a plan's feature, while it previously was assigned to goals. Goals are now instantiated by the *desired* clinical state attribute. Since plans selection is the primary focus of this study, we considered this design more efficient. Plans are in fact now filtered according to the clinical state to which they apply, the initial clinical state (ICS), and the current goal (G). This design also provides a mechanism for updating goals. When a plan has been successfully executed, its FCS is checked for equality with Goal's DCS. If they are equal, the desired clinical state has been achieved and the goal is deactivated. The implementation of this mechanism will be part of our future efforts. Two other attributes of the previous definition, the 'intentional verb' and the 'target', have been collapsed into the desired clinical state attribute. We deemed this as a more agile approach serving better our purpose of demonstrating the overall reasoning capabilities of the system. The utility of using verb/object, or verb/target, constructs in goal definition, is related to the possibility of controlling complex plans cycles. If, for instance, a plan to decrease blood pressure is activated, its stopping condition will be the achievement of normal blood pressure values. The goal's verb component 'decrease' can be used as a parameter to infer such condition. These types of inferences were not in the scope of this study and we opted for a more compact representation.

To handle more complex features than just provision of atomic plans, the current architecture would require richer plans' definitions. Temporal patterns for plans execution, such as time intervals, for instance, are not addressed here

Specification of goals by clinicians is also a potential limit in the usability of this system. The clinician in fact would have to specify his intention in the current design. The DCS attribute can be used to sort those plans with matching FCS and present to the clinician a list of pre-generated options from which they can select their immediate goals.

Protégé-OWL proved to be a sound editing and browsing tool, but some of its limitations would affect applications developed in real clinical settings. The lack of a multi user interface to support collaborative ontology development and maintenance might hinder sharing formalized knowledge over a medical community. Recently available commercial tools might offer support to build large-scale knowledge bases containing millions of concepts.

The evaluation of the proposed model was limited in several respects including the number of cases, the non-real setting, and the lack of unbiased judges. At this point of the project, it would have been premature to conduct a more thorough evaluation. While the results from the limited evaluation indicate that this approach is promising, these findings need to be validated in a real clinical setting.

Conclusion

This preliminary study presents a novel architecture for a medical decision-support system using ontological design and formal semantics for goals. These technologies proved to be effective in two principal ways: 1) in easily creating, managing and updating formally represented medical knowledge bases, and 2) in executing articulate queries on these knowledge bases in a distributed fashion.

The study also demonstrated the important role that intentions play in medical decision-making and how they can be exploited to improve the proficiency of automated decision-support systems.

References

1. Sackett D. L., Rosenberg W. M. C., Gray J. A. M., Haynes R. B., Richardson W. S. Evidence based medicine: what it is and what it isn't. BMJ 1996;312:71-72 (13 January).

2. Winikoff M., Padgham L., Harland J., Thangarajah J. Declarative & Procedural Goals in Intelligent Agent Systems. http://www.cs.rmit.edu.au/agents/Papers/KR2002.pdf.

3. Braubach L., Pokahr A., Moldt D., Lamersdorf W. Goal Representation for BDI Agent. Systems Lecture Notes in Computer Science; Springer-Verlag GmbH Volume 3346 / 2005, Chapter: p. 44.

4. Bratman M., Two faces of intention. Philosophical Review 93: 375–4051984.

5. Georgeff M., Pell B., Pollack M. The Belief-Desire-Intention Model of Agency Editors: Tambe M., Wooldridge M.

6. Noy N.F., Crubezy M., Fergerson R.W., Knublauch H., Tu S.W., Vendetti J., Musen M.A.. Protege-2000: an open-source ontology-development and knowledge-acquisition environment. AMIA Annu Symp Proc. 2003; 953.

7. Knublauch H., Musen M.A., Rector A.L. Editing Description Logic Ontologies with the Protégé OWL Plugin. International Workshop on Description Logics - DL2004, Whistler, BC, Canada (2004).

8. World Wide Web Consortium. OWL Web Ontology Language Semantics and Abstract Syntax. W3C Recommendation, 10 Feb, 2004.

9. Shahar Y., Miksch S., Johnson P. The Asgaard project: a task-specific framework for the application and critiquing of time-oriented clinical guidelines. Artif Intell Med. 1998 Sep-Oct;14(1-2):29-51.

10. Fox J., Johns N., Lyons C., Rahmanzadeh A., Thomson R., Wilson. PROforma: a general technology for clinical decision support systems. Comput Methods Programs Biomed. 1997 Sep;54(1-2):59-67.

11. Boxwala A.A., Peleg M., Tu S., Ogunyemi O., Zeng Q.T., Wang D., Patel V.L., Greenes R.A., Shortliffe E.H..GLIF3: a representation format for sharable computer-interpretable clinical practice guidelines. J Biomed Inform. 2004 Jun;37(3):147-61.

12. Rumbaugh J., Jacobson I., Booch G. The Unified Modeling Language reference manual. Addison-Wesley Longman Ltd., Essex, UK.

13. Institute for Clinical Systems Improvement. Diagnosis of Breast Disease. November 2003.

14. Institute for Clinical Systems Improvement, Hypertension Diagnosis and Treatment. February 2004.

15. Erol K., Hendler J., Nau D.S. Semantics for Hierarchical, Task-Network Planning Technical report CS-TR-3239, UMIACS-TR-94-31.

16. Haarslev V., Moller R. Racer: A Core Inference Engine for the Semantic Web http://www.sts.tu-harburg.de/~r.f.moeller/racer/papers/2003/HaMo03e.pdf.

17. Boyd C. M. Clinical Practice Guidelines and Quality of Care for Older Patients With Multiple Comorbid Diseases Implications for Pay for Performance, JAMA, August 10, 2005—Vol 294, No. 6.

18. Hashmi N., Boxwala A., Zaccagnini D., Fox J. Formal Representation of Medical Goals for Medical Guidelines. Proceedings, Medinfo 2003.