Mapping clinical guidelines representation primitives to decision theory concepts

Stefania Montani and **Paolo Terenziani**[1]

Clinical guidelines (GL) can be defined as a means for specifying the "best" clinical procedures and for standardizing them. In recent years, the medical community has started to recognize that a computer-based treatment of GL provides relevant advantages, such as automatic connection to the patient databases and, more interestingly, decision making facilities; thus, many different approaches and projects have been developed to this hand (see e.g. [4, 1]).

As a matter of fact, decision making is a central issue in clinical practice. In particular, supporting therapy selection is a critical objective to be achieved. Consider that, when executing a GL on a given patient, a physician can be faced with a choice among different therapeutic alternatives, and identifying the most suitable one is often not straightforward. Actually in several situations no alternative is really "better" than the others, from a strictly clinical viewpoint, and GL (unlike protocols) are only meant to present all the range of choices, leaving to the user the responsibility of selecting the "right" one. Even when resorting to a computer-based system for GL management, just "local" information, describing the decision at hand, are normally shown to the user. On the other hand, the possibility of obtaining a complete scenario of the decision consequences (in terms of the probability of the different therapy outcomes, of therapy utilities, and of money, time and resources spent following the different paths), would be clearly an added value for physicians. Moreover, also hospital administrators could take advantage of such a facility: by providing a projection of the (economic and temporal) cost of each alternative, it would allow for a quicker analysis and optimization of the patients care processes.

In clinical practice, various selection parameters (such as the costs and effectiveness of the different procedures) are sometimes available when executing a GL, but the task of comparing and balancing them is typically left to the physician. A system able to automatize the comparison and to provide quantitative results would be of great help in several real world situations.

Decision theory seems a natural candidate as a methodology for affording this analysis; nevertheless, rather interestingly, none of the systems for computerized GL management found in the literature embeds a decision theory tool able to compare therapeutic alternatives from the point of view of utilities and costs. A reason for this lack, in our opinion, is due to a difficulty in mapping the representation primitives between the two areas.

The contribution we provide in this paper is a knowledge representation one, aimed exactly at realizing this mapping in a general and reusable way.

The GL representation primitives adopted by the systems described in the literature may differ for several details, but if we look at them from at a more abstract level, we can identify a few skeletal concepts, shared by most of them. First, a GL can be represented as a graph, where nodes are the *actions* to be executed, and arcs are the *control relations* linking them. We can distinguish between *atomic* and *composite* actions (plans), where atomic actions represent simple steps in a guideline, and plans represent actions which can be defined in terms of their components via the *has-part* relation. The guideline itself is a plan. Three different types of *atomic actions* can then be identified (for the terminology used here, please refer in particular to [2, 8]): (1) *work actions*, i.e. actions that describe a procedure which must be executed at a given point of the guideline; (2) *query actions*, i.e. requests of information (typically patient's parameters), that can be obtained from the outside world (physicians, databases, patient's visits or interviews); (3) *decision actions*, used to model the selection among different alternatives. Decision actions can be further subdivided into *diagnostic decisions*, used to make explicit the identification of the disease the patient is suffering from, and *therapeutic decisions*, used to represent the choice of a path in the GL, containing the implementation of a particular therapeutic process. In this case, the choice will prune other paths, which implement different therapies, not as suitable as the selected one. The selected therapeutic process will typically be composed by several work actions and/or plans.

Actions in a GL are connected through *control* relations, which establish which actions can be executed next, and in what order. In particular the *alternative* relation describes how alternative paths can stem from a decision action, and the *repetition* relation states that an action has to be repeated several times (maybe a number of times not known a priori, until a certain exit condition becomes true).

In a well-formed GL, a decision action is preceded by a query action, that is adopted to collect all the patient's parameters necessary (and sufficient) for taking the decision itself (here we refer in particular to therapeutic decisions: in the GL context, as a matter of fact, a diagnostic decision only allows to classify the disease the patient is suffering from, and simply preludes to a therapeutic decision among suitable alternatives to care the patient herself). Each decision is therefore based on an (explicit or implicit) data collection completed at decision time, and does not depend on the previous history of the patient (i.e. on previous data collections and on previous decisions found along the path that leads to the decision at hand). We can thus say that the GL describes a first-order Markov model, since each time a query action is implemented, the patient's situation is completely re-assessed, and an (explicit or implicit) query action is always found before a decision action. This observation justifies the mapping of GL primitives to the field of decision theory, and in particular allows us to represent a GL as a Markov Decision Process (MDP), which has been recognized as a basic representation framework for dynamic decision making under uncertainty. Note that

[1] DI, Università del Piemonte Orientale "A. Avogadro", Spalto Marengo 33, 15100 Alessandria, Italy, e-mail: stefania,terenz@mfn.unipmn.it

Markov models have been widely used in the last decades in medical decision making, and represent nowadays a well-understood instrument to cope with time-dependent medical decision problems [6].

In particular, it is straightforward to define the concept of *state* as the set of patient's parameters that are normally measured for taking decisions and for assessing therapy outcomes. As already observed, the extraction of patient's parameters is obtained through a query action in the GL domain. Each parameter is a *state variable*. The query action is the means for assessing the patient's state at time t, where t is also the time instant at which the decision has to be taken.

On the other hand, understanding what GL primitives can produce state transitions requires a detailed analysis of the primitives themselves. It is well known that state transitions are changes in the patient's state variables, due to the effect of an action; deciding what action to take in order to maximize utility is the goal of decision theory. We can observe that: (1) query actions just photograph the state, and are not able to produce any change in it; (2) decision actions are used to make explicit which path will be followed, in a range of alternatives: in particular, diagnostic decisions only cover a disease classification task, while therapeutic decisions indicate what decision process has to be implemented, being a decision process a set of work actions (i.e. a plan); (3) work actions finally represent operative steps in the guideline, such as providing a drug. Therefore, each work action will potentially have an effect on the patient's state variables. Thus, we will consider work actions as the means to produce state transitions.

From the analysis above, it is possible to model the GL process as a discrete-time one, where time discretization is performed by the query actions preceding therapeutic decisions. Given this discretization, the process is also completely observable (in a GL, a decision can be taken only if all the required parameters have been collected: if some needed data are missing, the query action will wait for them and the decision will be delayed). The state transition between time t (time of the first therapeutic decision) and time $t+1$ (time of the following one) is due to *all* the work actions between the two time instants. Therefore, there is no quantitative mapping between the times of state assessments and the chronological dates at which the actions of the GL take place. A set of work actions can take days as well as months to be completed, but state assessment will always take place at the following query action; so, the actual temporal distance between t and $t+1$ can be very dishomogeneous from case to case.

The *utility* of reaching a state can be evaluated in terms of life expectancy, corrected by Quality Adjusted Life Year (QALYs) [3]. We can derive the utility of a state from the medical literature, as we do for obtaining the probability of state transitions. Note that, for those medical fields in which the medical literature does not provide these numbers, it is reasonable to expect this information to be available in the near future. As a matter of fact, the increasing exploitation of Hospital Information Systems and of computerized GL management tools will allow for the collection of large amounts of clinical practice data, on which it will be easy to draw statistics, at least at the local level. Consider also that relying on local data is not necessarily a limitation: actually a guideline always needs to be contextualized to the features of the hospital in which it has to be implemented, before its exploitation begins [9].

The discussion about *costs*, on the other hand, is more complex. Costs can be interpreted as monetary expenses: each work action (e.g. buying a drug, or using diagnostic instrumentation) will typically have a price. But costs can also be evaluated in terms of the time and the resources required to complete work actions. Whatever is the unit chosen to quantify costs, costs are not a property of the state reached after a transition, but depend on the work actions that have to be implemented along the selected path.

As a final remark, note that we have identified work actions as the only possible responsible for changes in the patient's parameters (i.e. in the state variables). As a matter of fact, this is a simplification, since patient's parameters can vary due to *exogenous* reasons (e.g. because the patient becomes older, or because she catches another disease). To explicitly represent all these possibilities, we should be provided with a whole model of the patient's behavior and of all the stochastic variables that could influence it. This kind of information is normally not explicitly available in clinical practice. Even though such a model was available, on the other hand, it would not provide relevant advantages for decision support: exogenous effects could change the patient's state, but they cannot be controlled by the physician; therefore their explicit representation is useless for supporting therapeutic decisions, and the state transition model can be simplified. Moreover, the effect of exogenous factors is implicitly taken into account by deriving transition probability values from the medical literature.

As anticipated, our knowledge representation contribution could be resorted to by any of the GL systems in the literature, since the basic primitives we treat are shared by most of them. In particular, even though some approaches do not distinguish between diagnostic decisions and therapeutic ones, and do not explicit the query action before each therapeutic decision, these concepts can always be found in the systems' underlying semantics. We believe that a decision theory tool would provide a valuable support to physicians, thus reinforcing the claim that the adoption of AI techniques can lead to relevant advantages in the (semi)-automatic treatment of clinical guidelines, favoring the dissemination and the actual adoption of computer science tools within the medical community. A first step in this direction is being taken by the system GLARE, where a facility based on the analysis described in this work is being implemented [8, 5]; in GLARE, we plan to represent the MDP describing the GL resorting to a *dynamic decision network* [7], a choice that allows one to explicitly take advantage of conditional independencies from the modeling viewpoint, and to rely on several powerful algorithms for probabilistic inference.

REFERENCES

[1] D.B. Fridsma (Guest Ed.). Special issue on workflow management and clinical guidelines. *Journal of the American Medical Informatics Association*, 1(22):1-80, 2001.

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

[3] M.R. Gold, J. E. Siegel, L. B. Russell, M. C. Weinstein. *Cost-Effectiveness in Health and Medicine*. Oxford University Press, New York, 1996.

[4] C. Gordon and J.P. Christensen, editors. *Health Telematics for Clinical Guidelines and Protocols*. IOS Press, Amsterdam, 1995.

[5] S. Montani, P.Terenziani. Decision theory issues for supporting therapy selection in computerized guidelines. Accepted for IPMU 2004, Perugia, July 2004.

[6] F.A. Sonnenberg, J.R. Beck. Markov models in medical decision making:a practical guide. *Medical Decision Making*, 13:332-338, 1993.

[7] J. Tatman, R. Shachter. Dynamic programming and influence diagrams. *IEEE Transactions on Systems, Man and Cybernetics*, 20:365-379, 1990.

[8] P. Terenziani, G. Molino, M. Torchio. A modular approach for representing and executing clinical guidelines. *Artificial Intelligence in Medicine*, 23(3):249-276, 2001.

[9] P. Terenziani, S. Montani, A. Bottrighi, M. Torchio, G. Molino, G. Correndo. A context-adaptable approach to clinical guidelines. Accepted for MedInfo 2004, San Francisco, September 2004.