

Toward Probabilistic Analysis of Guidelines

Arjen Hommersom

Radboud University Nijmegen
Institute for Computing and Information Sciences
Nijmegen, The Netherlands
arjenh@cs.ru.nl *

Abstract. In the formal analysis of health-care, there is little work that combines probabilistic and temporal reasoning. On the one hand, there are those that aim to support the clinical thinking process, which is characterised by trade-off decision making taking into account uncertainty and preferences, i.e., the process has a probabilistic and decision-theoretic flavour. On the other hand, the management of care, e.g., guidelines and planning of tasks, is typically modelled symbolically using temporal, non-probabilistic, methods. This paper proposes a new framework for combining temporal reasoning with probabilistic decision making. The framework is instantiated with a guideline modelling language combined with probabilistic pharmacokinetics and applied to the treatment of diabetes mellitus type 2.

1 Introduction

Clinical guidelines are highly structured documents providing appropriate standards of care. Recommendations of clinical guidelines are based around actions, also sometimes referred to as tasks or interventions, that physicians are advised to perform when treating specific groups of patients. If there is only one possible action in the care pathway, the situation is relatively easy. For many diseases, however, there are at least several actions that have to be performed. In order to model such care pathways, expressive formalisms were developed such as Asbru, *PROforma*, and GLARE (see [2]). These languages are best characterised by the term ‘task-network models’, as they model the guideline as a network of component tasks that unfold over time [11].

In order to ensure quality of these guidelines, several symbolic analysis approaches have been proposed, such as simulation and verification, e.g., using formal methods for checking that resulting guidelines comply to certain quality criteria (see, e.g., [7]). Whereas the recommendations can be looked upon as symbolic and logical entities, the underlying knowledge, typically medical evidence, involves uncertainty. So far, this aspect of a guideline has been largely ignored, but is required for a deeper analysis of the quality of guidelines.

* Visiting researcher at the Department of Computer Science, Katholieke Universiteit Leuven, Celestijnenlaan 200A, BE-3001 Heverlee, Belgium.

The representation of uncertainty is particularly important if one wants to support the guideline development process. As far as we are aware, there are no knowledge representation formalisms specifically designed for this task. Indeed, the main challenge for such a representation is to integrate knowledge underlying guidelines with medical decision making taking into account uncertainty derived from the evidence. A complicating factor is that the current process of developing recommendations from uncertain knowledge is hardly clear. For example, the ‘*The guidelines manual*’ by the British NHS [10] states the following about recommendations: “If (the guideline) combines consideration of several possible interventions, it may include discussion of the position of an intervention within a pathway of care”. There is no way to decide *how* to determine the position of the intervention within a pathway of care. In this paper, a language is introduced for exploring different possible combinations of treatments and to consider their outcomes. This also opens up a new possibility for *personalisation* of medicine as different treatment options can be explored for specific patients.

This paper is organised as follows. In Section 2, we will introduce the probabilistic framework. The implementation of this framework is discussed in Section 3 and applied to the management of diabetes mellitus type 2 in Section 4. Finally, in Section 5, we conclude and discuss future work.

2 Probabilistic Clinical Model

In this section, the probabilistic model underlying our analysis is introduced. First, we discuss the model from a symbolic point of view, after which it is refined with probabilistic aspects.

2.1 Histories

As a point of departure, we look at care as a process that can be modelled as sequences of possible actions and conditions, that we previously called *histories*. An elaborate treatment can be found in [6], where histories were studied from a more general and theoretical point of view. In short, histories describe sequences of the state of a patient and interventions being performed on that patient. Expectations extend histories with possible states in the future, constrained by clinical management and (patho-)physiological processes. Formally, we define a history H as a sequence of triples $\{(p_k, i_k, t_k)\}_{k=0}^{k=n}$, where p_k is a (patient) state, i_k a description of interventions, and t_k a time-point from a partially ordered set. We denote s_k for the *state* consisting of the combination of the patient state and the performed interventions, i.e., $s_k = (p_k, i_k)$. In this paper, the complete state is modelled as a set of attribute-value pairs $\langle a_i, v_i \rangle$, where a_i is an attribute and v_i the value of that attribute.

The set of all histories is called \mathcal{H} . An expectation E extends a history to other possible histories, i.e., it is a function $E : \mathcal{H} \rightarrow \mathcal{P}(\mathcal{H})$, where $\mathcal{P}(\mathcal{H})$ denotes the power set of \mathcal{H} , together with some boundary condition (see [6, Chapter

7]) ensuring that these expectations are sound. Sound expectations are step-functions E_s , for which holds that if $H' \in E_s(H)$, for some H that contains n triples, then $H' = H \cup \{(p_{n+1}, i_{n+1}, t_{n+1})\}$, where $t_{n+1} > t_n$, i.e., E_s extends H with information at a later time-point and at no other time-point.

2.2 Probabilistic histories

The intent of expectations is to describe possible continuations of the history. Not all of these continuations are equally likely, however. In order to model this uncertainty, we define a *probabilistic expectation* as a function E_p which associates each possible expectation with a probability, i.e.,

$$E_p : \mathcal{H} \rightarrow \mathcal{P}(\mathcal{H} \times [0, 1])$$

such that E_p is a step function, i.e., it extends H with just one new time-point in the future. Furthermore, the set of possible expectations is complete, i.e., if $E_p(H) = \{(h_1, P_1), \dots, (h_n, P_n)\}$, then $\sum P_i = 1$.

In the probabilistic model, we assume (i) time-invariance, i.e., the probability of the expectation only depends on states, but not on the times and (ii) the expectations are *Markovian*, that is, expectations depend on the present (the last state), but not on the past. Note that there are no restrictions to prevent the embedding of all information about the past in the last state of the history, except that it will impact the complexity of reasoning. If these two assumptions are combined, expectations only depend on the final state (s_n) of a history, thus expectation can be described by a transition relation:

$$s_n \xrightarrow{P(s_n, s_{n+1})} s_{n+1}$$

which yields a new history $h_{n+1} = \{\dots, (s_n, t_n), (s_{n+1}, t_{n+1})\}$, with $t_{n+1} > t_n$, as one of the expectations of the history h_n .

To further refine the model, the transition relation is then decomposed according to the state decomposition, i.e., the transition of going from v to v' is given for all attribute-value $\langle a, v \rangle$ pairs such that:

$$s_n, \langle a, v \rangle \xrightarrow{P(c(a, v, v', s_n))} \langle a, v' \rangle$$

where $c(a, v, v', s_n)$ models a random (choice) variable, which is assumed to be independent of choices for other attributes in state s_n . The complete transition probability is defined as the conjunction of each of the choices. From this, and the fact that the choices are independent of each other, we obtain:

$$P(s_n, s_{n+1}) = P\left(\bigwedge_{i=0}^m c(a_i, v_i, v'_i, s_n)\right) = \prod_{i=1}^m P(c(a_i, v_i, v'_i, s_n))$$

where $s_n = \{\langle a_i, v_i \rangle\}_{i=0}^m$ and $s_{n+1} = \{\langle a_i, v'_i \rangle\}_{i=0}^m$.

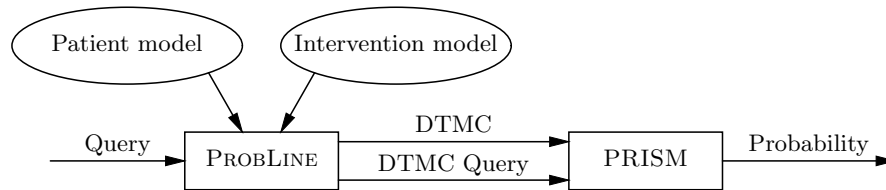


Fig. 1. General framework of PROBLINE. From a patient and intervention model it generates a discrete-time Markov chain (DTMC), which is used to compute the answer to a query.

3 ProbLine

3.1 General framework

Probabilistic histories act as a framework for a software tool tool called PROBLINE. Fig. 1 provides an overview of the approach. PROBLINE provides an interface for a patient model as well as an intervention model and answers probabilistic queries in the form $P(a_i = v_i | t)$, where $\langle a_i, v_i \rangle \in p_i$ and t a point in time, i.e., the probability that a patient attribute a_i has value v_i at time t . PROBLINE, as presented in this paper, is fully implemented and runs on the YAP system, which is a recent high-performance implementation of Prolog.

The intervention model and patient model describe probabilistic transitions as formally discussed in the probabilistic clinical model. As the system makes no other assumptions, different knowledge representation formalisms could be used to construct intervention and patient models.

Probabilistic histories can be interpreted as infinite discrete-time Markov chains (DTMC), i.e., a DTMC with a potentially infinite number of states. As all probabilistic queries are given a time t , and given that there are a finite number of transitions from a given state, the query can be computed on the basis of a finite DTMC. The computation of this finite DTMC that can answer the query is handled by the tool, which then calls the probabilistic symbolic model checking tool PRISM [9] to compute the answer to the query. This model checker incorporates state-of-the art symbolic data structures and algorithms for computing probabilistic queries of complex models.

3.2 Patient model

As far as we are aware, there are no dedicated knowledge representations for patient models. We therefore use the following basic representation. First, we can declare new patient attributes as follows:

- `patient_attribute(+A, +V)`
Declares a new attribute A with an initial value V .
- `patient_attribute(+A, +V, +P)`
Declares a new attribute A with possible initial values V (i.e., V is a list) and a list with probabilities P with a initial probability for each value in V .

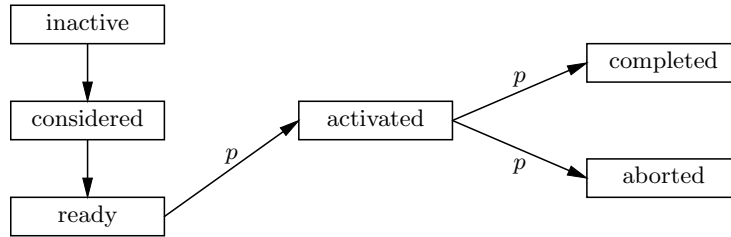


Fig. 2. Simplified Asbru state chart of [1]. The state transitions that might be probabilistic are indicated by p . The remaining two transitions are technical in the sense that they are independent of any user interaction. Conditions to go from one state to the other are not shown in this figure.

The transition system can then be modelled using the following two primitives:

- `choice(+A, +V, +V', +S, +P)`
This corresponds to the choice operator as discussed in the previous section, where A is an attribute, V is a value, V' is an updated value, S is the state, and P is the probability of making this choice.
- `patient_av(+S, +A, -V)`
Provides the value V of an attribute A in state S .

A graphical representation, for example a flowchart, might be more appropriate in practice, and can be built on top of these predicates.

3.3 Intervention model

As said, PROBLINE is not restricted to a specific knowledge representation. Nevertheless, in order to illustrate its capabilities, we instantiate the intervention model with a computer-interpretable guideline modelling language. In this work, we have chosen a small subset of Asbru and modelled its state-chart semantics [1]. A simplified version of the state chart semantics of Asbru is given in Fig. 2. In the original semantics there are non-deterministic choices to go from one state to the other, e.g., whether or not to abort a task if the abort condition holds. In this model, we include the possibility of probabilistic transitions between states. In particular, the transition from ready to activated, and from activated to some terminated state, i.e., completed or aborted state, is a probabilistic transition. These probabilities model the chance that a physician acts if a treatment is allowed to start or could be terminated.

The representation of tasks, in Asbru they are called plans, is implemented with the following four predicates:

- `plan_body(+N, +T, +W, +C)`
Defines a new plan with name N , with a body type T (currently either ‘sequential’ or ‘parallel’). This plan has a list of children C and has a wait-for condition W in order to model optional and mandatory plans for N .

- `{abort,complete,filter}_condition(+N, +S)`:
Specifies in which state S the abort/complete/filter condition is true (see [1]).
If it is not specified for a state S , then it is ‘false’ by default because of the usual Prolog semantics (negation as failure). These conditions influence the state transitions in the Asbru semantics, e.g., a plan can only abort if the `abort_condition` holds. Note that arbitrary Prolog programs can be used to specify when such a condition holds.

The Asbru semantics as mentioned is implemented as a module of PROBLINE and could be extended with other features of the language, e.g., retry conditions and time-annotations.

4 Diabetes mellitus type 2

In the previous sections, the probabilistic framework of PROBLINE, including its syntax and semantics, have been discussed. In this section, we apply this framework to the management of diabetes mellitus type 2.

4.1 Management of the disease

It is well known that diabetes type 2 is a complicated disease: various metabolic control mechanisms are deranged and many different organ systems may be affected by the disorder. Patho-physiologically, there are two main phenomena, namely, insufficient secretion of the hormone insulin due to a decreased production of insulin by *B cells* in the Langerhans islets of the *pancreas*, and insulin resistance in liver, muscle and fat tissue. For the individual patient, there is a lot of uncertainty to which extent these phenomena occur, which makes it difficult to predict whether an intervention will be effective. Not only is the underlying physiology uncertain, choosing which drug to give a patient is also based on other considerations such as availability, cost, safety, tolerability, and convenience. Personalised medicine promises a path for individually optimised treatment choices, but realising this promise will require a more comprehensive characterisation of disease and drug response.

For diabetes, quite a lot of known about the effect of drugs. In this paper, we will focus on a well-known biguanide drug called metformin, which is commonly prescribed as the primary oral anti-diabetic. The dosage that we will consider is 2,000 mg/day, which is for many patients optimal [13]. The efficacy of metformin 2,000 mg/day has also been estimated [3]: the fasting plasma glucose (FPG) will be lowered up to 86 mg/dL +/- 10 mg/dL (95% confidence interval). The steady-state situation is not reached before 8 days and a linear reduction of the FPG seems appropriate [8]. Recently, Shu et al. [14] integrated diverse data supporting the hypothesis that genetic variations in the encoding of a protein called organic cation transporter 1 (OCT1) affects the response to metformin. These variants of OCT1 lead to half of the bio-availability of metformin. It was also estimated that about 20% of the Caucasian population carries one of these mutations. No cases are known with multiple mutations, nor is it common for other ethnicities.

```

plan_body(metformin, sequential, [], []).
abort_condition(metformin,S) :-
    patient_av(S, glycemia, hyper),
    patient_av(S, metformin_app, T),
    T >= 10.
complete_condition(metformin,S) :- patient_av(S, glycemia, normo).
filter_condition(metformin,_).

plan_body(treatment, sequential, [metformin], [metformin]).
abort_condition(treatment, _) :- false.
complete_condition(treatment,_).
filter_condition(treatment,_).

```

Fig. 3. Model of metformin application in Asbru. The treatment will only be aborted if the time $T \geq 10$, where the time granularity is in days.

4.2 Probabilistic model of metformin pharmacokinetics

Given the information above, we estimate a probability distribution for a variable `max_reduction` such that it is a discretised normal variable with mean 86 (mg/dL) and standard deviation 5 as 95% of the population are within 2 standard deviations (i.e., 10 mg/dL) for a normal distribution. Of course, if the raw data is available, then other distributions could be used instead. Here we are limited to the published mean and 95% confidence interval, for which we assume that a bell-shaped distribution is appropriate. Then, to describe expected probabilities of reaching normoglycemia, we have deterministic transitions:

```

choice(glycemia, hyper, normo, S, 1) :-
    patient_av(S, baselineFPG, FGL),
    patient_av(S, time_metformin_application, T),
    patient_av(S, oct1_variant, Oct),
    patient_av(S, max_reduction, Max),
    expected_normo(FGL, AdT, Max, Oct).

```

where `FPG` is a FPG at baseline, `T` is the time that metformin is applied, and `Oct` is a binary variable that is true if the patient has a variation of OCT1 affecting the efficacy of metformin. The predicate `expected_normo` computes the expected FPG based on these parameters and returns true if the expected FPG is less than 110 mg/dL, which is commonly defined as normoglycemia.

4.3 Experiments

Using PROBLINE, we can now answer a number of question surrounding the treatment with metformin. We provide a model of metformin application in Asbru in Fig. 3 where transition probabilities are set to 0.5 (see Fig. 2).

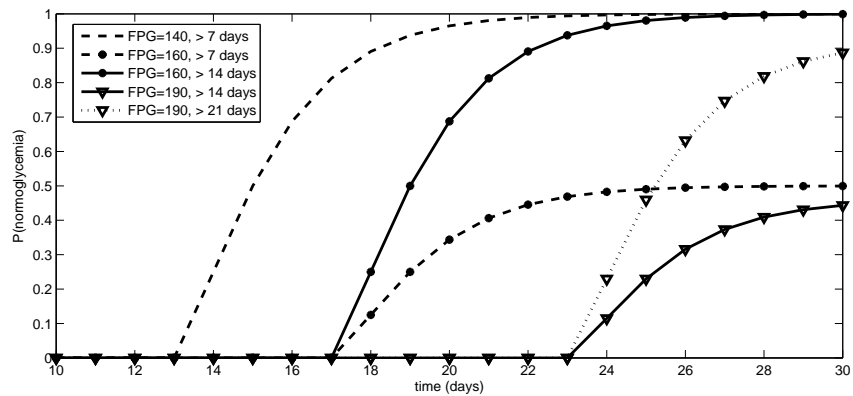


Fig. 4. Probabilistic simulation of metformin application to patients with different FPG at baseline. Time of metformin application is varied as well.

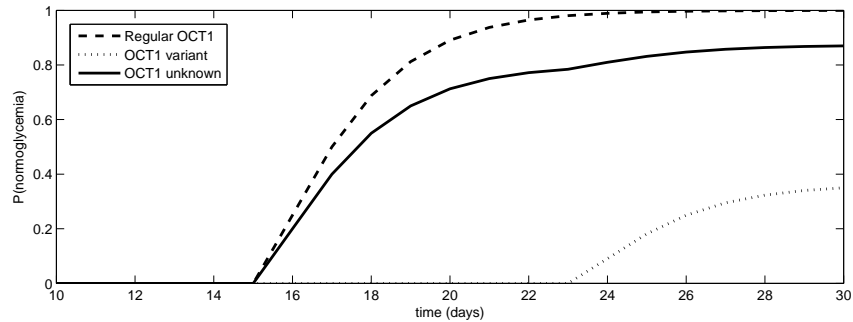


Fig. 5. Probabilistic simulation of metformin application to patients with or without a variation in the OCT1 protein.

Question 1: How long should metformin be applied before it can be decided to stop the treatment? There is a trade-off here: if the treatment with oral anti-diabetics is stopped too early then patients may be injecting themselves with insulin for no good reason; if the treatment is stopped too late, then patients who need treatment with insulin are not treated appropriately. In Fig. 4, we plot a number of dose-response curves for different patients (without OCT1 variations). For people with an initial low fasting plasma glucose, the effect of treatment is relatively quick, whereas people with an initial high fasting plasma glucose, the effect is much slower and might not be effective at all even after prolonged treatment.

Question 2: What improvement could we gain using genetic information? As mentioned, it is hypothesised that the OCT1 protein plays a key role in the

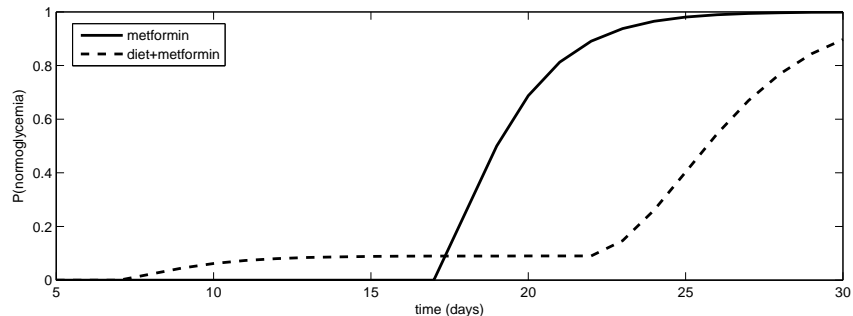


Fig. 6. Comparing metformin with and without diet in patients with normal OCT1.

efficacy of metformin. This suggests that it would be useful to test whether a patient has a variation in this gene before treatment. In Fig. 5, two patients are plotted with the same FPG at baseline but with different OCT1 proteins. On average, patients in this population (baseline FPG=150) have a good chance that metformin is effective. However, for the patients with the OCT1-variant, the chance that metformin is effective is rather small and it might be better to prescribe an alternative drug. In the end, such pharmacogenetics could be used for the personalisation of treatments [12].

Question 3: Should we try diet before metformin? In the management of diabetes, one of the first things that is recommended is a diet. However, the efficacy is low [5]. If we include this information, then we can ask PROBLINE to compare a treatment with or without a diet. Fig. 6 shows the difference between the two, where it seems that diet is of little benefit. Of course, there could be other reasons to recommend a diet, such as improved health or less side-effects of the treatment. The point is that PROBLINE allows for the exploration of different treatments taking into account uncertainty.

5 Conclusions

In this paper, we presented a new method for the analysis of care processes taking into account uncertainty. The system implementing this theory, PROBLINE, can handle typical task-network representations of guidelines and flowchart-like patient models. We presented a case-study in the treatment of diabetes mellitus type 2 illustrating the strength of this approach.

This work can be extended in several ways. In this paper, we introduced the core of the language consisting of reasoning with the dynamic aspects of guidelines, taking into account uncertainty. In future work, we would like to extend the language with other types of knowledge derived from, for example, medical ontologies and Bayesian networks. To accomplish this, a more powerful probabilistic logical language is required, e.g., using one of the recently developed logics in the field of statistical relational learning (see [4]).

References

1. M. Balsler, C. Duelli, and W. Reif. Formal semantics of Asbru - an overview. In *Proceedings of the International Conference on Integrated Design and Process Technology*, Pasadena, June 2002. Society for Design and Process Science.
2. P. de Clercq, K. Kaiser, and A. Hasman. Computer-interpretable guideline formalisms. In A ten Teije, S. Miksch, and P.J.F. Lucas, editors, *Computer-based Medical Guidelines and Protocols: A Primer and Current Trends*, pages 22–43. IOS Press, 2008.
3. A.J. Garber, T.G. Duncan, A.M. Goodman, D.J. Mills, and J.L. Rohlf. Efficacy of metformin in type ii diabetes: results of a double-blind, placebo-controlled, dose-response trial. *Am J Med*, 103(6):491–507, 1997.
4. L. Getoor and B. Taskar, editors. *Introduction to Statistical Relational Learning*. Adaptive Computation and Machine Learning. MIT Press, 2007.
5. L.S. Hermann, B. Scherstén, P.O. Bitzén, T. Kjellström, F. Lindgärde, and A. Melander. Therapeutic comparison of metformin and sulfonylurea, alone and in various combinations. *Diabetes Care*, 16(10):1100–1109, 1994.
6. A.J. Hommersom. *On the Application of Formal Methods to Clinical Guidelines*. PhD thesis, University of Nijmegen, 2008.
7. A.J. Hommersom, P.C. Groot, M. Balsler, and P.J.F. Lucas. Formal methods for verification of clinical practice guidelines. In A ten Teije, S. Miksch, and P.J.F. Lucas, editors, *Computer-based Medical Guidelines and Protocols: A Primer and Current Trends*, pages 63–80. IOS Press, 2008.
8. Y. Hong, S. Rohatagi, Habtemariam. B., J.R. Walker, S.L. Schwartz, and D.E. Mager. Population exposure-response modeling of metformin inpatients with type 2 diabetes mellitus. *J Clin Pharmacol*, 48(6):696–707, 2008.
9. M. Kwiatkowska, G. Norman, and D. Parker. PRISM: Probabilistic symbolic model checker. In T. Field, P. Harrison, J. Bradley, and U Harder, editors, *Proceedings TOOLS'02*, volume 2324 of *LNCS*, pages 200–204. Springer, 2002.
10. NHS. The guidelines manual. http://www.nice.org.uk/media/68D/3C/The_guidelines_manual_2009_-_Chapter_9_Developing_and_wording_guideline_recommendations.pdf.
11. M. Peleg, S. Tu, J. Bury, P. Ciccarese, and J. Fox. Comparing computer-interpretable guideline models: a case-study approach. *Journal of the American Medical Informatics Association*, 10(1):52–68, 2003.
12. M.L. Reitman and E.E. Schadt. Pharmacogenetics of metformin response: a step in the path toward personalized medicine. *J Clin Invest*, 117(5), 2007.
13. J.H.B. Scarpello and H.C.S. Howlett. Metformin therapy and clinical uses. *Diab Vasc Dis Res*, 5(3):157–167, 2008.
14. Y. Shu, S.A. Sheardown, C. Brown, R.P. Owen, S. Zhang, R.A. Castro, A.G. Ianculescu, L. Yue, J.C. Lo, E.G. Burchard, C.M. Brett, and K.M. Giacomini. Effect of genetic variation in the organic cation transporter 1 (OCT1) on metformin action. *J Clin Invest*, 117:1422–1431, 2007.