— Submitted to the 8th World Congress on Medical Informatics (Medinfo ’95) —

OPTIMIZING DIAGNOSTIC AND THERAPEUTIC STRATEGIES USING DECISION-THEORETIC PLANNING: PRINCIPLES AND APPLICATIONS

Charles E. Kahn, Jr., M.D.[a] and Peter Haddawy, Ph.D.[b]

[a] Section of Information and Decision Sciences, Department of Radiology, Medical College of Wisconsin, Milwaukee, Wisconsin 53226, USA (*ckahn@mcw.edu*).

[b] Department of Electrical Engineering and Computer Science, University of Wisconsin–Milwaukee, Milwaukee, Wisconsin 53201, USA (*haddawy@cs.uwm.edu*).

ABSTRACT

Objective: Decision-theoretic planning is a new technique for selecting optimal actions. The authors sought to determine whether decision-theoretic planning could be applied to medical decision making to identify optimal strategies for diagnosis and therapy.

Methods: An existing model of acute deep venous thrombosis (DVT) of the lower extremities — in which 24 management strategies were compared — was converted into a set of conditional-probabilistic actions for use by the DRIPS decision-theoretic planning system. Actions were grouped into an abstraction/decomposition hierarchy. A utility function was defined in accordance with the existing DVT management model to incorporate the costs and risks of the diagnostic tests and treatments.

Results: From 18 primitive actions (such as “perform venography” and “treat if venography shows thigh DVT”), a total of 312 possible concrete plans were encoded within the abstraction/decomposition hierarchy. The DRIPS planning system used abstraction techniques to eliminate 136 possible plans (44%) from consideration. It determined that, given the parameters specified, the most cost-effective management strategy was “no tests, no treatment.” This result differed from the published result of “perform ultrasonography, treat if positive.” In reviewing the original article, it was determined that DRIPS had revealed an error in the manually constructed decision trees used in that manuscript. At values for the cost of death of $75,000 and greater, the optimal strategy became “impedance plethysmography (IPG), don’t wait, perform venography if IPG is positive, and treat only if venography shows thigh DVT.”

Conclusion: Decision-theoretic planning is applicable to medical decision making, and may be an extremely useful technique for complex decisions. The use of inheritance abstraction makes the technique computationally tractable for complex planning problems, and the modular nature of the data entry may help eliminate errors that appear in manually encoded decision trees.

1. INTRODUCTION

Decision trees are used with increasing frequency to help physicians select patient management strategies that offer the most cost-effective tests and treatments. The models needed to create these decision trees may be complex, and researchers often must develop a large number of decision trees to evaluate competing strategies. Constructing these decision trees can be a laborious, time-consuming process. To simplify this process, we investigated the application of decision-theoretic refinement planning to medical decision making.

Decision-theoretic refinement planning is a new planning approach that seeks to find the plan that maximizes expected utility relative to a user-defined utility function. It reasons with actions organized into an abstraction hierarchy; this hierarchy allows the planner to eliminate classes of plans without evaluating them individually. We applied the DRIPS decision-theoretic refinement planning system [1] to data from a published analysis of cost-effective management of acute lower-extremity deep venous thrombosis, in which 24 decision trees were evaluated [2]. Abstract actions (*e.g.*, test) and their sub-actions (*e.g.*, venography) were encoded into a hierarchy, and the effects of each action were defined in terms of conditional probabilities. A utility function was defined based on the specified costs and risks. This article demonstrates the use of DRIPS to identify optimal medical management strategies.

2. DECISION-THEORETIC PLANNING

The DRIPS decision-theoretic refinement planning system provides a mechanism to identify optimal plans involving diagnosis and treatment [1]. DRIPS embodies a probabilistic model of the world, incorporates temporal reasoning, and allows users to define the utility of various possible outcomes. The DRIPS system's goal is to find the optimal plan — the one that offers the greatest expected utility. In all but the simplest of domains, analysis of each possible plan by exhaustive enumeration would be computationally prohibitive. DRIPS provides a method for abstracting probabilistic actions; it can reason with these abstractions in such a way that suboptimal classes of plans can be eliminated without explicitly examining all plans in those classes.

A planning problem is described in terms of a probabilistic model of the world, an initial state distribution, a set of available actions, and a utility function. The world is represented as a set of attributes (variables) with associated probability distributions. Attributes may be discrete or continuous. We assume that changes to the world are limited to those effects explicitly described in the action descriptions, so we do not allow for exogenous events. DRIPS uses a goal-directed utility model for deadline and maintenance goals [3].

A plan is a sequence of actions. Actions are both conditional and probabilistic: under certain conditions an action will have a given effect with a given conditional probability. Actions that change the state of the world (*e.g.*,, treatments) and actions involving observations (*e.g.*, diagnostic tests) are treated in a uniform manner. This allows one to include information such as the risks of diagnostic procedures (*e.g.*, the risk of venography causing DVT).

Inheritance abstraction formalizes the notion that analogous actions can be grouped together into an action class characterized by common features [4]. DRIPS searches through a space of plans structured into an abstraction/decomposition hierarchy. An abstract action has one or more sub-actions, which themselves may be abstractions or concrete actions. A decomposible action has one or more sub-plans that all must be executed in sequence.

DRIPS finds the optimal plan by building abstract plans, comparing them, and refining only those that might yield the optimal plan. Unlike "generative" planning schemes that build plans by adding actions, DRIPS performs "refinement" planning. It begins with an abstract plan, and subsequently refines the plan from more general to more specific. An abstract plan's outcomes are sets of outcomes of more concrete plans. Since different probability and utility values may be associated with each specific outcome, in general a probability range and a utility range will be associated with each abstract outcome.

The expected utility of an abstract plan is represented by an interval, which includes the expected utilities of all possible instantiations of that abstract plan. Refining the plan, i.e., instantiating one of its actions, tends to narrow the interval. When the expected utility intervals of two plans do not overlap, the one with the lower interval can be eliminated. (Or, if the upper bound on the expected utility of each

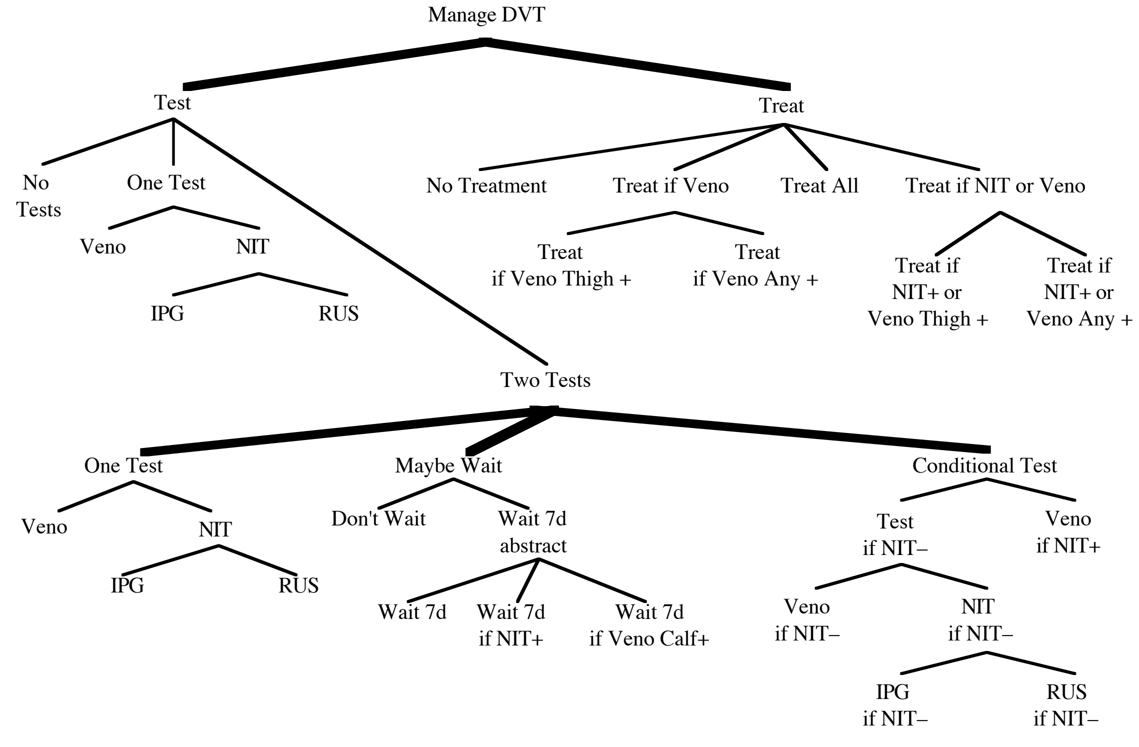

Figure 1. Abstraction-decomposition diagram of actions for DVT decision model. Thick lines represent decompositions; thin lines represent abstractions.

abstract plan is tight, i.e., there is an instance of the abstract plan with that expected utility, then at each refinement step we can eliminate all abstract plans but the one with the highest upper bound.) Its algorithm is guaranteed to find the optimal plan [5]; with an appropriate abstraction hierarchy, its complexity is exponentially better than that of exhaustive enumeration [1].

The outcome of a plan is a probability distribution over a set of chronicles. We need to compute this distribution from the initial state distribution and the action descriptions. We make the Markovian assumption that the conditional probabilities of an action's effects given the action and the conditions in its description are independent of all other conditions at the same time or earlier and all other previous actions. We also assume that the conditions determining the effects of an action are probabilistically independent of the action. Under these assumptions, we can compute the outcome distribution of an action by multiplying the conditional probabilities describing the action effects by the probabilities of the conditions. Projecting a plan produces a probability distribution over a future-branching tree of chronicles.

3. ACUTE DEEP VENOUS THROMBOSIS

Appropriate management of patients with suspected acute deep venous thrombosis (DVT) of the lower extremities remains an important and complex clinical problem. The clinical findings of DVT do not permit diagnosis with certainty [6]. Unchecked, lower-extremity DVT can progress to pulmonary embolism, a condition that entails significant morbidity and mortality. Anticoagulation therapy for DVT is expensive and carries the risk of severe hemorrhage. Even diagnostic procedures such as venography entail risks such as contrast reaction and iatrogenic DVT.

To evaluate the applicability of DRIPS to medical decision making, we constructed a model for diagnosis and treatment of DVT. The model was based on data from a manuscript that compared 24 different management strategies [2]. We chose this manuscript because it contained explicit probability and utility data, and because we hoped to show the advantage of our technique over the manual construction of decision trees. Our model did not include all possible combinations of actions; specifically, we limited the model to two tests, with a possible seven-day period between them. The test procedures included contrast venography (Veno) and two non-invasive tests (NITs): impedance plethysmography (IPG) and real-time ultrasonography (RUS). Treatment, which consisted solely of anticoagulation therapy, included unconditional actions (*e.g.*, Treat All) and conditional actions (*e.g.*, treat if thigh DVT seen on venography [Treat if Veno Thigh+]). The actions were organized into an abstraction-decomposition hierarchy (**Figure 1**).

In this hierarchy, the Manage DVT action decomposes into two actions, Test and Treat, which are performed in sequence. The Test action is an abstraction for three possible sub-actions: No Tests, One Test, or Two Tests. The results of each action, whether concrete or abstract, are described using conditional probabilities (**Figure 2**). Our model for management of suspected DVT encompasses 312 concrete plans; for example, one complete plan is "IPG, Wait 7d if NIT-, Veno if NIT-, Treat if Veno Any+."

4. RESULTS

Based on the "standard model" assumptions in the model by Hillner and colleagues [2], in which the "cost" of death was set at $30,000, DRIPS determined that the plan with the highest expected utility was "No Tests, No Treatment." This plan differed from the one identified in their manuscript, which was "Treat if RUS positive, Observe if RUS negative." In fact, for several plans, DRIPS calculated expected utility values that differed from those published in the source manuscript. In attempting to resolve this discrepancy, we discovered that at least one of the 24 decision trees constructed by Hillner and colleagues was incorrect (Hillner BE, personal communication).

In evaluating this model, DRIPS used its abstraction hierarchy to evaluate only 176 plans instead of the universe of 312 plans. In another domain, DRIPS has achieved a "pruning factor" in excess of 96 percent [7]. Such pruning is important to gain computational efficiency in large models. Analysis of the DVT management model required 8 minutes on a DEC 5000/240; the software is written in Common Lisp, and due to its developmental nature, has not been optimized.

We studied the effect of the "cost" of death on the selection of the optimal management plan. At values of $30,000 (the "standard model" value in the source manuscript [2], $40,000, and $60,000, the optimal

```
(add-action NIT                                                             [1]
            (thigh-dvt NITResult cost dead pe)                              [2]
            (cond ((and (thigh-dvt = 1) (dead = 0) (pe = 0))
                   (1 (NITResult = 2) 0.95
                      (NITResult = 1) 0.05
                      (cost = cost + (120 300)) 1))
                  ((and (thigh-dvt = 0) (dead = 0) (pe = 0))                [3]
                   (1 (NITResult = 1) (0.90 0.98)
                      (NITResult = 2) (0.02 0.10)
                      (cost = cost + (120 300)) 1))
                  ((or (dead = 1) (pe = 1)) (1 )))
            (list IPG RUS))                                                 [4]
```

Figure 2. The description of the abstract action NIT (non-invasive test) specifies the action title [1] and the variables that affect and are affected by the action [2]. A conditional action [3] is described: here, if thigh DVT is absent, the patient is not dead, and there is no pulmonary embolism (PE), then the test result is negative (NITResult=1) with an interval probability of .90 to .98 or positive (NITResult=2) with an interval probability of .02 to .10. The cost, which is part of the utility function, is increased from 120 to 300. Line [4] lists the abstracted actions, namely IPG (impedance plethysmography) and RUS (real-time ultrasonography). Action descriptions are written in a pseudo-Lisp macro language.

plan was “No Tests, No Treatment.” At values of $75,000, $100,000, and $200,000 for the cost of death, the optimal plan became “impedance plethysmography (IPG), Don’t Wait, perform venography if IPG was positive (Veno if NIT+), and treat only if venography shows thigh DVT (Treat if Veno Thigh+).”

5. DISCUSSION

Decision-theoretic planning, as embodied in the DRIPS planning system, is applicable to medical decision making and may be very useful for a wide variety of medical problems. The current investigation has shown its applicability to problems of protocol development, but the planner also could be applied to the management of individual patients. Strengths of this approach include the use of probabilistic reasoning, the incorporation of utility functions, and the ability to reason about time, and the ability to reason with both discrete and continuous variables. DRIPS incorporates diagnostic procedures, their costs, and their potential side-effects; few planning systems are able to deal with diagnostic actions.

The use of DRIPS offers several practical advantages as well. DRIPS allows the model builder to provide “modules” of information: the description of each action can be highly compartmentalized. The abstraction hierarchy allows one to think qualitatively about the model to make certain that all possible actions are included. One this hierarchy is defined, the system performs the rote task of evaluating all possible combinations of actions. As was demonstrated in our DVT model, the use of DRIPS is less tedious and less prone to error than construction of multiple decision trees. One can easily alter a model’s underlying assumptions to evaluate the effects on the outcome; in this way, DRIPS can perform sensitivity analysis of the model.

New tests and treatment can be added easily to DRIPS planning models. We intend to complete the DVT model by adding branches for three or four possible tests before considering treatment, which will result in a universe of more than 100,000 concrete plans for DRIPS to evaluate;. Clearly, the use of abstraction will be important in this model. We also are working to improve the DRIPS user interface and to automatically generate abstractions from the descriptions of “primitive” concrete actions. We view decision-theoretic planning as a technique of wide applicability to medical decision making.

REFERENCES

1. Haddawy P, Suwandi M. Decision-theoretic refinement planning using inheritance abstraction. In: Hammond K, ed. Proceedings of the Second International Conference on Artificial Intelligence Planning Systems. Menlo Park, CA: AAAI Press, 1994: 266-271.
2. Hillner BE, Philbrick JT, Becker DM. Optimal management of suspected lower-extremity deep vein thrombosis: an evaluation with cost assessment of 24 management strategies. Arch Intern Med 1992;152:165-175.
3. Haddawy P, Hanks S. Representations for decision-theoretic planning: utility functions for deadline goals. In: Nebel B, Rich C, Swartout W, ed. Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference. San Mateo, CA: Morgan Kaufmann, 1992: 71-82.
4. Tenenberg JD. Reasoning about Plans. San Mateo, CA: Morgan Kaufmann, 1991:213-283.
5. Haddawy P, Doan A. Abstracting probabilistic actions. Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence. San Mateo, CA: Morgan Kaufmann, 1994: (in press).
6. Landefeld S, McGuire E, Cohen AM. Clinical findings associated with acute proximal deep vein thrombosis: a basis for qualifying clinical judgment. Am J Med 1990;88:382-388.
7. Haddawy P, Bach D. BREWER: using decision-theoretic refinement planning to select optimal beer brewing processes. (In preparation).