Influence diagrams for medical decision problems: Some limitations and proposed solutions*

Manuel Luque and **Francisco Javier Diez**
Dept. Inteligencia Artificial. UNED
28040 Madrid. Spain
mluque@bec.uned.es, fjdiez@dia.uned.es

Carlos Disdier
Pulmonary Section
San Pedro Alcantara Hospital
1004 Caceres. Spain
cdisdier@separ.es

Abstract

When trying to solve two medical decision problems we have encountered several difficulties: how to represent and operate with decomposable utility functions, how to calibrate our human experts and explain them the "reasoning" of our influence diagrams, and how to deal with partially ordered decisions. This paper describes these difficulties and the solutions we have adopted.

1 Introduction

One of the medical problems we are currently working on is the mediastinal staging of non-small cell lung cancer. There are several tests available, such as computed tomography scan (CT scan), transbronchial needle aspiration (TBNA), mediastinoscopy (MED) and others, which have different characteristics of sensitivity, specificity, morbidity and mortality. The other problem on which we are working is the management of mild head injury.

Influence diagrams are a framework which serves as an effective modeling tool for decision problems. An influence diagram (ID) [3], consists of a directed acyclic graph having three kinds of nodes: decision (graphically represented by squares or rectangles), chance (circles or ovals), and utilities (diamonds). Each decision node represents to actions under the direct control of the decision maker. Each chance node represents a random variable. In medical IDs, utility nodes represent medical outcomes and costs (morbidity, mortality, economic cost...).

The quantitative information that defines an ID is given by assigning to each chance node X_i a probability distribution $p(X_i|pa(X_i))$, where $pa(X_i)$ represents the parents of the node X_i in the graph, and assigning to the utility node U a function $\psi(pa(U))$. The objective of the evaluation of an influence diagram is obtaining a policy for each decision, which prescribes a set of optimal actions for the decision maker. The policy for each decision is a function of the variables that are known when the decision is made.

*This research has been supported by the Spanish Ministry of Science and Technology under grant TIC-2001-2973-C05-04.

2 Limitations of influence diagrams for medical decision problems

This section describes the difficulties we have found when building those IDs and how we have extended Elvira to cope with them. Elvira[1] is a Java tool to construct probabilistic decision support systems. Elvira works with Bayesian networks (BN) and influence diagrams and it can operate with discrete and continuous variables. It has an easy Graphical User Interface (GUI) for constructing BNs and IDs.

2.1 Decomposable utilities

An essential component of an influence diagram is the utility function. In its original formulation [3], each ID had only one utility node, which entails several disadvantages. First, the human expert has to assess more parameters. Second, the bigger the utility function the more time and memory space is required for the computational evaluation of the ID. Third, policies tend to include more variables than when using decomposable utility functions.

In order to overcome these shortcomings, Tatman and Shachter [5] introduced a new kind of utility node, called super-value nodes, which represent a function of their parents' utilities, and proposed an algorithm for evaluating such generalized IDs.

The first limitation we encountered when building our medical IDs is that current software tools do not admit super-value nodes. At most, they accept several utility nodes under the assumption that the global utility is the sum of all of them. For this reason, we extended Elvira's GUI and format so that it could cope with super-value nodes. Figure 1 shows the current version of our ID for mediastinal staging. The three rectangles represent the decisions: one of them represents the decision of performing a TBNA or not, the second represents the decision about performing a mediastinoscopy, and the third represents the decision about the treatment, which can be thoracotomy, radiotherapy, chemotherapy, or palliative care. Rounded rectangles represent chance variables: one is the main diagnosis (N2-N3), three are tests, and the forth indicates whether the patient survives the mediastinoscopy. At the bottom, there are seven utility nodes; two of them are super-value nodes, which indicates a decomposition of the utility function.

[1] See `http://www.ia.uned.es/~elvira` and [1].

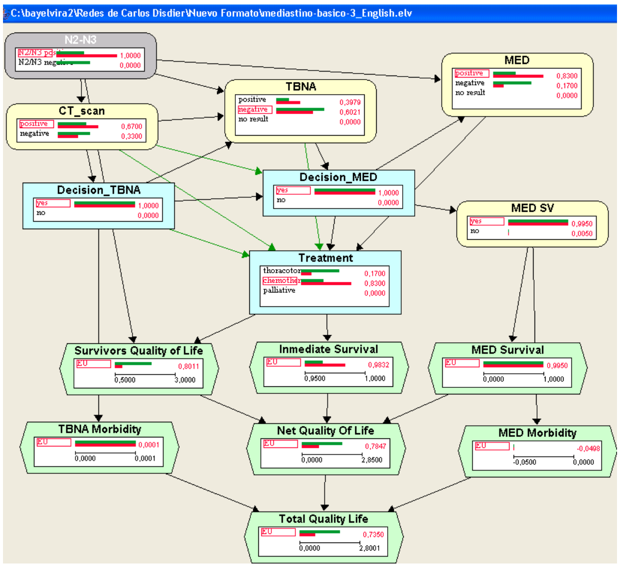

Figure 1: Explanation for a medical decision problem in Elvira

Furthermore, we realized that the algorithm of Tatman and Shachter was unsatisfactory because it is based on arc reversal, an operation that involves inefficient divisions of potentials and often introduces unnecessary variables in the resulting policies. For this reason, we developed a new algorithm which is in general more efficient and does not tend to introduce so many unnecessary variables in the policies [4].

2.2 Explanation in influence diagrams

One of the key factors for the acceptance of expert systems in real world domains and especially in medicine is the capability to explain their reasoning. For this reason we have extended Elvira's explanation facilities from Bayesian networks to IDs. In addition to showing the resulting policies by opening a window for each decision, Elvira can also display numerical and graphical information inside each node (see Figure 1): horizontal bars inside the nodes represent the probability that a chance variable takes on a certain value, the probability that the decision maker chooses a certain action for a decision,[2] or the expected utility for a utility node.

In Elvira it is possible to assign values to chance and decision variables in the same way as evidence is assigned to the corresponding nodes of a Bayesian network. It is also possible to show several "evidence cases", i.e., the probabilities and utilities for several subpopulations. For instance, Figure 1 shows two horizontal bars for each value of a chance or decision variable and for each utility node, thus comparing the situation in which we know the patient belongs to the N2-N3 positive group with the situation in which we have not any previous information. Please note that the node *Total Expected Utility* shows the global utility, while other utility nodes, such as *Survivors Quality of Life*, represent partial utilities.

[2]Policies are deterministic functions, but it makes sense to speak of the probability of an action because decisions are based on chance variables.

2.3 Order of decisions

The traditional definition of the influence diagrams assume that there is an order between decisions. However, in some medical problems the question is just which tests should be performed and in what order. For example, in the mediastinal staging problem, we may wonder what is the best order among the three tests, CT scan, MED and TBNA.

Several representations of decision problems have been proposed in order to let have a partial order between decisions (see [2] and references therein). The task of finding the best order is performed by the evaluation algorithm. The aim of these representations is representing and solving asymmetric decision problems. However, they tend to obscure the structure of the medical decision problems, because they are too general and do not consider the specific characteristics of the medical decision problems. This makes more difficult any kind of explanation of the reasoning and less efficient the evaluation of the decision problem.

For this reason we are currently exploring new representational schemes that will lead to more simple and intuitive influence diagrams, and in turn would require new algorithms for their evaluation.

3 Conclusion

In this paper we have discussed the shortcomings of influence diagrams for solving real-world problems in medicine, as well as the limitations of some of the current algorithms and software packages. We have also shown some of our current efforts aimed at having more efficient representation schemes, algorithms and software tools.

References

[1] The Elvira Consortium. Elvira: An environment for creating and using probabilistic graphical models. In *Proceedings of the First European Workshop on Probabilistic Graphical Models (PGM'02)*, pages 1–11, Cuenca, Spain, 2002.

[2] T. D. Nielsen F. V. Jensen and P. Shenoy. Sequential influence diagrams: A unified asymmetry framework. In *Proceedings of the Second European Workshop on Probabilistic Graphical Models*, pages 121–128. P. Lucas (eds.), 2004.

[3] R. A. Howard and J. E. Matheson. Influence diagrams. In R. A. Howard and J. E. Matheson, editors, *Readings on the Principles and Applications of Decision Analysis*, pages 719–762. Strategic Decisions Group, Menlo Park, CA, 1984.

[4] M. Luque and F. J. Diez. Variable elimination for influence diagrams with super-value nodes. In *Proceedings of the Second European Workshop on Probabilistic Graphical Models*, pages 145–152. P. Lucas (eds.), 2004.

[5] J. A. Tatman and R. D. Shachter. Dynamic programming and influence diagrams. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(2):365–379, 1990.