Decision analysis with influence diagrams using Elvira's explanation facilities

Manuel Luque and Francisco J. Díez
Dept. Inteligencia Artificial, UNED
Juan del Rosal, 16, 28040 Madrid, Spain

Abstract

Explanation of reasoning in expert systems is necessary for debugging the knowledge base, for facilitating their acceptance by human users, and for using them as tutoring systems. Influence diagrams have proved to be effective tools for building decision-support systems, but explanation of their reasoning is difficult, because inference in probabilistic graphical models seems to have little relation with human thinking. The current paper describes some explanation capabilities for influence diagrams and how they have been implemented in Elvira, a public software tool.

1 Introduction

Influence diagrams (IDs) are a probabilistic graphical model for modelling decision problems. They constitute a simple graphical formalism that makes it possible to represent the three main components of decision problems: the uncertainty, due to the existence of variables not controlled by the decision maker, the decisions or actions under their direct control, and their preferences about the possible outcomes.

In the context of expert systems, either probabilistic or heuristic, the development of explanation facilities is important for three main reasons (Lacave and Díez, 2002). First, because the construction of those systems with the help of human experts is a difficult and time-consuming task, prone to errors and omissions. An explanation tool can help the experts and the knowledge engineers to *debug* the system when it does not yield the expected results and even before a malfunction occurs. Second, because human beings are reluctant to *accept* the advice offered by a machine if they are not able to understand how the system arrived at those recommendations; this reluctancy is especially clear in medicine (Wallis and Shortliffe, 1984). And third, because an expert system used as an intelligent *tutor* must be able to communicate to the apprentice the knowledge it contains, the way in which the knowledge has been applied for arriving at a conclusion, and what would have happened if the user had introduced different pieces of evidence (what-if reasoning).

These reasons are especially relevant in the case of probabilistic expert systems, because the elicitation of probabilities is a difficult task that usually requires debugging and refinement, and because the algorithms for the computation of probabilities and utilities are, at least apparently, very different from human reasoning.

Unfortunately, most expert systems and commercial tools available today, either heuristic or probabilistic, have no explanation capabilities. In this paper we describe some explanation methods developed as a response to the needs that we have detected when building and debugging medical expert systems (Díez et al., 1997; Luque et al., 2005) and when teaching probabilistic graphical models to pre- and postgraduate students of computer science and medicine (Díez, 2004). These new methods have been implemented in Elvira, a public software tool.

1.1 Elvira

Elvira[1] is a tool for building and evaluating graphical probabilistic models (Elvira Consor-

[1] At http://www.ia.uned.es/∼elvira it is possible to obtain the source code and several technical documents about Elvira.

tium, 2002) developed as a join project of several Spanish universities. It contains a graphical interface for editing networks, with specific options for canonical models, exact and approximate algorithms for both discrete and continuous variables, explanation facilities, learning methods for building networks from databases, etc. Although some of the algorithms can work with both discrete and continuous variables, most of the explanation capabilities assume that all the variables are discrete.

2 Influence diagrams

2.1 Definition of an ID

An influence diagram (ID) consists of a directed acyclic graph that contains three kinds of nodes: *chance nodes* $\mathbf{V}_C$, *decision nodes* $\mathbf{V}_D$ and *utility nodes* $\mathbf{V}_U$—see Figure 1. Chance nodes represent random variables not controlled by the decision maker. Decision nodes correspond to actions under the direct control of the decision maker. Utility nodes represent the decision maker's preferences. Utility nodes can not be parents of chance or decision nodes. Given that each node represents a variable, we will use the terms variable and node interchangeably.

In the extended framework proposed by Tatman and Shachter (1990) there are two kinds of utility nodes: *ordinary utility nodes*, whose parents are decision and/or chance nodes, and *super-value nodes*, whose parents are utility nodes, and can be in turn of two types, sum and product. We assume that there is a utility node U_0, which is either the only utility node or a descendant of all the other utility nodes, and therefore has no children.[2]

There are three kinds of arcs in an ID, depending on the type of node they go into. Arcs into chance nodes represent probabilistic dependency. Arcs into decision nodes represent availability of information, i.e., an arc $X \to D$ means that the state of X is known when making decision D. Arcs into utility nodes represent functional dependence: for ordinary utility nodes, they represent the domain of the associated utility function; for a super-value node they indicate that the associated utility is a function (sum or product) of the utility functions of its parents.

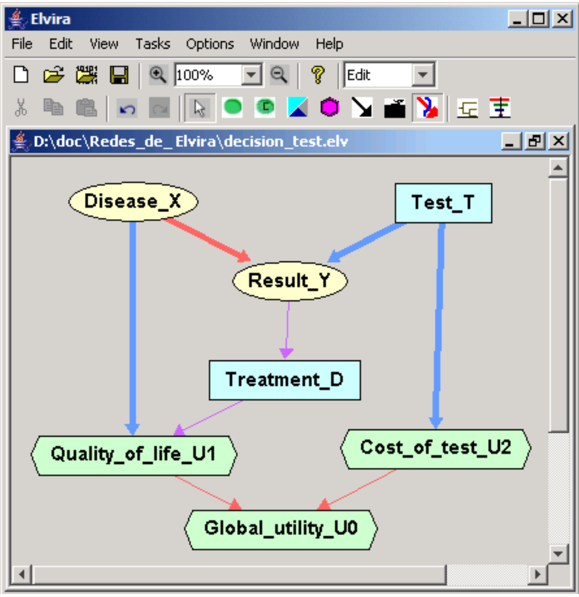

Figure 1: ID with two decisions (ovals), two chance nodes (squares) and three utility nodes (diamonds). There is a directed path including all the decisions and node U_0.

Standard IDs require that there is a directed path that includes all the decision nodes and indicates the order in which the decisions are made. This in turn induces a partition of $\mathbf{V}_C$ such that for an ID having n decisions $\{D_0, \ldots, D_{n-1}\}$, the partition contains $n+1$ subsets $\{\mathbf{C}_0, \mathbf{C}_1, ..., \mathbf{C}_n\}$, where $\mathbf{C}_i$ is the set of chance variables C such that there is a link $C \to D_i$ and no link $C \to D_j$ with $j < i$; i.e., $\mathbf{C}_i$ represents the set of random variables known for D_i and unknown for previous decisions. $\mathbf{C}_n$ is the set of variables having no link to any decision, i.e., the variables whose true value is never known directly.

Given a chance or decision variable V, two decisions D_i and D_j such that $i < j$, and two links $V \to D_i$ and $V \to D_j$, the former link

[2]An ID that does not fulfill this condition can be transformed by adding a super-value node U_0 of type sum whose parents are the utility nodes that did not have descendants. The expected utility and the optimal strategy (both defined below) of the transformed diagram are the same as those of the original one.

is said to be a *no-forgetting link*. In the above example, $T \to D$ would be a non-forgetting link.

The variables known to the decision maker when deciding on D_i are called *informational predecessors* of D_i and denoted by $IPred(D_i)$. Assuming the *no-forgetting hypothesis*, we have that $IPred(D_i) = IPred(D_{i-1}) \cup \{D_{i-1}\} \cup \mathbf{C}_i = \mathbf{C}_0 \cup \{D_0\} \cup \mathbf{C}_1 \cup \ldots \cup \{D_{i-1}\} \cup \mathbf{C}_i$.

The quantitative information that defines an ID is given by assigning to each random node C a probability distribution $P(c|pa(C))$ for each configuration of its parents, $pa(C)$, assigning to each ordinary utility node U a function $\psi_U(pa(U))$ that maps each configuration of its parents onto a real number, and assigning a utility-combination function to each super-value node. The domain of each function U is given by its *functional predecessors*, $FPred(U)$. For an ordinary utility node, $FPred(U) = Pa(U)$, and for a super-value node $FPred(U) = \bigcup_{U' \in Pa(U)} FPred(U')$.

For instance, in the ID of Figure 1, we have that $FPred(U_1) = \{X, D\}$, $FPred(U_2) = \{T\}$ and $FPred(U_0) = \{X, D, T\}$.

In order to simplify our notation, we will sometimes assume without loss of generality that for any utility node U we have that $FPred(U) = \mathbf{V}_C \cup \mathbf{V}_D$.

2.2 Policies and expected utilities

For each configuration $\mathbf{v}_D$ of the decision variables $\mathbf{V}_D$ we have a joint distribution over the set of random variables $\mathbf{V}_C$:

$$P(\mathbf{v}_C : \mathbf{v}_D) = \prod_{C \in \mathbf{V}_C} P(c|pa(C)) \qquad (1)$$

which represents the probability of configuration $\mathbf{v}_C$ when the decision variables are externally set to the values given by $\mathbf{v}_D$ (Cowell et al., 1999).

A *stochastic policy* for a decision D is a probability distribution defined over D and conditioned on the set of its informational predecessors, $P_D(d|IPred(D))$. If P_D is degenerate (consisting of ones and zeros only) then we say the policy is deterministic.

A *strategy* Δ for an ID is a set of policies, one for each decision, $\{P_D | D \in \mathbf{V}_D\}$. A strategy Δ induces a joint distribution over $\mathbf{V}_C \cup \mathbf{V}_D$ defined by

$$\begin{aligned} &P_\Delta(\mathbf{v}_C, \mathbf{v}_D) \\ &\quad = P(\mathbf{v}_C : \mathbf{v}_D) \prod_{D \in \mathbf{V}_D} P_D(d|IPred(D)) \\ &\quad = \prod_{C \in \mathbf{V}_C} P(c|pa(C)) \prod_{D \in \mathbf{V}_D} P_D(d|pa(D)) \quad (2) \end{aligned}$$

Let I be an ID, Δ a strategy for I and $\mathbf{r}$ a configuration defined over a set of variables $\mathbf{R} \subseteq \mathbf{V}_C \cup \mathbf{V}_D$ such that $P_\Delta(\mathbf{r}) \neq 0$. The probability distribution *induced by strategy* Δ *given the configuration* $\mathbf{r}$, defined over $\mathbf{R}' = (\mathbf{V}_C \cup \mathbf{V}_D) \setminus \mathbf{R}$, is given by:

$$P_\Delta(\mathbf{r}'|\mathbf{r}) = \frac{P_\Delta(\mathbf{r}, \mathbf{r}')}{P_\Delta(\mathbf{r})} \ . \qquad (3)$$

Using this distribution we can compute the *expected utility of* U *under strategy* Δ *given the configuration* $\mathbf{r}$ as:

$$EU_U(\Delta, \mathbf{r}) = \sum_{\mathbf{r}'} P_\Delta(\mathbf{r}'|\mathbf{r}) \psi_U(\mathbf{r}, \mathbf{r}') \ . \qquad (4)$$

For the terminal utility node U_0, $EU_{U_0}(\Delta, \mathbf{r})$ is said to be the *expected utility of strategy* Δ *given the configuration* $\mathbf{r}$, and denoted by $EU(\Delta, \mathbf{r})$.

We define the *expected utility of* U *under strategy* Δ as $EU_U(\Delta) = EU_U(\Delta, \blacklozenge)$, where $\blacklozenge$ is the empty configuration. We also define the *expected utility of strategy* Δ as $EU(\Delta) = EU_{U_0}(\Delta)$. We have that

$$EU_U(\Delta) = \sum_{\mathbf{r}} P_\Delta(\mathbf{r}) EU_U(\Delta, \mathbf{r}) \ . \qquad (5)$$

An *optimal strategy* is a strategy Δ_{opt} that maximizes the expected utility:

$$\Delta_{opt} = \arg\max_{\Delta \in \Delta^*} EU(\Delta) \ , \qquad (6)$$

where Δ^* is the set of all strategies for I. Each policy in an optimal strategy is said to be an *optimal policy*. The *maximum expected utility* (MEU) is

$$MEU = EU(\Delta_{opt}) = \max_{\Delta \in \Delta^*} EU(\Delta) \ . \qquad (7)$$

The evaluation of an ID consists in finding the *MEU* and an optimal strategy. It can be proved (Cowell et al., 1999; Jensen, 2001) that

$$MEU = \sum_{\mathbf{c}_0} \max_{d_0} \ldots \sum_{\mathbf{c}_{n-1}} \max_{d_{n-1}} \sum_{\mathbf{c}_n} P(\mathbf{v}_C : \mathbf{v}_D)\psi(\mathbf{v}_C, \mathbf{v}_D) \; . \quad (8)$$

For instance, the *MEU* for the ID in Figure 1, assuming that U_0 is of type sum, is

$$MEU = \max_t \sum_y \max_d \sum_x P(x) \cdot P(y|t,x) \cdot \underbrace{(U_1(x,d) + U_2(t))}_{U_0(x,d,t)} \; . \quad (9)$$

2.3 Cooper policy networks

When a strategy $\Delta = \{P_D | D \in \mathbf{V}_D\}$ is defined for an ID, we can convert this into a Bayesian network, called *Cooper policy network* (CPN), as follows: each decision D is replaced by a chance node with probability potential P_D and parents $IPred(D)$, and each utility node U is converted into a chance node whose parents are its functional predecessors (cf. Sec. 2.1); the values of new chance variables are $\{+u, \neg u\}$ and its probability is $P_{CPN}(+u|FPred(U)) = norm_U(U(FPred(U)))$, where $norm_U$ is a bijective linear transformation that maps the utilities onto the interval $[0,1]$ (Cooper, 1988).

For instance, the CPN for the ID in Figure 1 is displayed in Figure 2. Please note the addition of the non-forgetting link $T \rightarrow D$ and that the parents of node U_0 are no longer U_1 and U_2 but T, X, and D, which were chance or decision nodes in the ID.

The joint distribution of the CPN is:

$$P_{CPN}(\mathbf{v}_C, \mathbf{v}_D, \mathbf{v}_U) = P_\Delta(\mathbf{v}_C, \mathbf{v}_D) \prod_{U \in \mathbf{V}_U} P_U(u|pa(U)) \; . \quad (10)$$

For a configuration $\mathbf{r}$ defined over a set of variables $\mathbf{R} \subseteq \mathbf{V}_C \cup \mathbf{V}_D$ and U a utility node, it is possible to prove that

$$P_{CPN}(\mathbf{r}) = P_\Delta(\mathbf{r}) \quad (11)$$

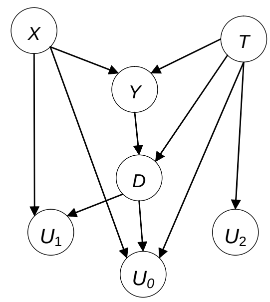

Figure 2: Cooper policy network (PN) for the ID in Figure 1.

$$P_{CPN}(+u|\mathbf{r}) = norm_U(EU_U(\Delta, \mathbf{r})) \; . \quad (12)$$

This equation allows us to compute $EU_U(\Delta, \mathbf{r})$ as $norm_U^{-1}(P_{CPN}(+u|\mathbf{r}))$; i.e., the expected utility for a node U in the ID can be computed from the marginal probability of the corresponding node in the CPN.

3 Explanation of influence diagrams in Elvira

3.1 Explanation of the model

The explanation of IDs in Elvira is based, to a great extent, on the methods developed for explanation of Bayesian networks (Lacave et al., 2000; Lacave et al., 2006b). One of the methods that have proven to be more useful is the automatic colorings of links. The definitions in (Lacave et al., 2006b) for the sign of influence and magnitude of influence, inspired on (Wellman, 1990), have been adapted to incoming arcs to ordinary utility nodes.

Specifically, the *magnitude of the influence* gives a relative measure of how a variable is influencing an ordinary utility node (see (Lacave et al., 2006a) for further details). Then, the influence of a link pointing to a utility node is positive when higher values of A lead to higher utilities. Cooper's transformation $norm_U$ guarantees that the magnitude of the influence is normalized.

For instance, in Figure 1 the link $X \rightarrow Y$ is colored in red because it represents a positive influence: the presence of the disease increases the probability of a positive result of the test. The link $X \rightarrow U_1$ is colored in blue because it represents a negative influence: the disease decreases the expected quality of life. The link $D \rightarrow U_1$ is colored in purple because the influence that it represents is undefined: the treatment is beneficial for patients suffering from X but detrimental for healthy patients.

Additionally, when a decision node has been assigned a policy, either by optimization or imposed by the user (see Sec. 3.3), the corresponding probability distribution P_D can be used to color the links pointing to that node, as shown in Figure 1.

The coloring of links has been very useful in Elvira for debugging Bayesian networks (Lacave et al., 2006b) and IDs, in particular for detecting some cases in which the numerical probabilities did not correspond to the qualitative relations determined by human experts, and also for checking the effect that the informational predecessors of a decision node have on its policy.

3.2 Explanation of reasoning: cases of evidence

In Section 2.3 we have seen that, given a strategy, an ID can be converted into a CPN, which is a true Bayesian network. Consequently, all the explanation capabilities for BNs, such as Elvira's ability to manage evidence cases are also available for IDs.

A *finding* states with certainty the value taken on a chance or decision variable. A set of findings is called *evidence* and corresponds to a certain configuration $\mathbf{e}$ of a set of observed variables $\mathbf{E}$. An *evidence case* is determined by an evidence $\mathbf{e}$ and the posterior probabilities and the expected utilities induced by it.

A distinguishable feature of Elvira is its ability to manage several evidence cases simultaneously. A special evidence case is the *prior case*, which is the first case created and corresponds to the absence of evidence.

One of the evidence cases is marked as the *current case*. Its probabilities and utilities are displayed in the nodes. For the rest of the cases, the probabilities and utilities are displayed only by bars, in order to visually compare how they vary when new evidence is introduced.

The information displayed for nodes depends on the kind of node. Chance and decision nodes present bars and numbers corresponding to posterior probabilities of their states, $P_\Delta(v|\mathbf{e})$, as given by Equation 3. This is the probability that a chance variable takes a certain value or the probability that the decision maker chooses a certain option (Nilsson and Jensen, 1998)—please note that in Equation 3 there is no distinction between chance and decision nodes. Utility nodes show the expected utilities, $EU_U(\Delta, \mathbf{e})$, given by Equation 4. The guide bar indicates the range of the utilities.

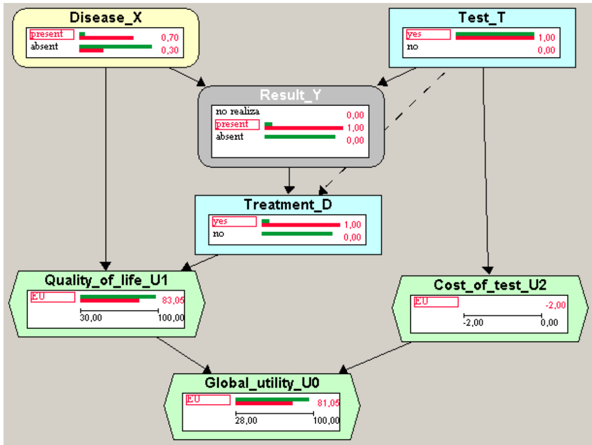

Figure 3: ID in Elvira with two evidence cases: (a) the prior case (no evidence); (b) a case with the evidence $\mathbf{e} = \{+y\}$.

Figure 3 shows the result of evaluating the ID in Figure 1. The link $T \rightarrow D$ is drawn as a discontinuous arrow to indicate that it has been added by the evaluation of the diagram. In this example, as there was no policy imposed by the user (see below), Elvira computed the optimal strategy. In this figure two evidence cases are displayed. The first one is the *prior case*, i.e., the case in which there is no evidence. The second evidence case is given by $\mathbf{e} = \{+y\}$; i.e., it displays the probabilities and utilities of the subpopulation of patients in which the test gives a positive result. Node Y is colored in gray to highlight the fact that there is evidence about it. The probability of $+x$, represented by a red bar, is 0.70; the green bar close to it represents the probability of $+x$ for the prior case, i.e., the prevalence of the disease; the red bar is longer than the green one because $P(+x|+y) > P(+x)$. The global utility for the second case is 81.05, which is smaller than the green bar close to it (the expected utility for the general population) because the presence of the symptom worsens the prognosis. The red bar for Treatment=yes is the probability that a patient having a positive result in the test receives the treatment; this probability is 100% because the optimal strategy determines that all symptomatic patients must be treated.

The possibility of introducing evidence in Elvira has been useful for building IDs in medicine: when we were interested in computing the posterior probability of diagnoses given several sets of findings, we need to manually convert the ID into a Bayesian network by removing decision and utility nodes, and each time the ID was modified we have to convert it into a Bayesian network to compute the probabilities. Now the probabilities can be computed directly on the ID.

3.2.1 Clarifying the concept of evidence in influence diagrams

In order to avoid confusions, we must mention that Ezawa's method (1998) for introducing evidence in IDs is very different from the way that is introduced in Elvira. For Ezawa, the introduction of evidence $\mathbf{e}$ leads to a different decision problem in which the values of the variables in $\mathbf{E}$ would be known with certainty before making any decision. For instance, introducing evidence $\{+x\}$ in the ID in Figure 1 would imply that X would be known when making decisions T and D. Therefore, the expected utility of the new decision problem would be

$$\max_t \sum_y \max_d P(y|+x:t,d) \cdot \underbrace{(U_1(+x,d) + U_2(t))}_{U_0(+x,d,t)} .$$

where $P(y|+x:t,d) = P(+x,y:t,d)/P(+x) = P(y|+x:t)$. In spite of the apparent similarity of this expression with Equation 9, the optimal strategy changes significantly to "always treat, without testing", because if we know with certainty that the disease X is present the result of the test is irrelevant. The MEU for this decision problem is $U_1(+x,+d)$.

In contrast, the introduction of evidence in Elvira (which may include "findings" for decision variables as well) does not lead to a new decision scenario nor to a different strategy, since the strategy is determined *before* introducing the "evidence". Put another way, in Elvira we adopt the point view of an external observer of a system that includes the decision maker as one of its components. The probabilities and expected utilities given by Equations 2 and 4 are those corresponding to the subpopulation indicated by $\mathbf{e}$ when treated with strategy Δ. For instance, given the evidence $\{+x\}$, the probability $P_\Delta(+t|+x)$ shown by Elvira is the probability that a patient suffering from X receives the test, which is 100% (it was 0% in Ezawa's scenario), and $P_\Delta(+d|+x)$ is the probability that he receives the treatment; contrary to Ezawa's scenario, this probability may differ from 100% because of false negatives. The expected utility for a patient suffering from X is

$$EU(\Delta, \{+x\}) = \\ = \sum_{t,y,d} P_\Delta(t,y,d|+x) \cdot \underbrace{(U_1(+x,d) + U_2(t))}_{U_0(+x,d,t)} .$$

where $P_\Delta(t,y,d|+x) = P_\Delta(t) \cdot P(y|t,+x) \cdot P_\Delta(d|t,y)$.

Finally, we must underlie that both approaches are not rivals. They correspond to different points of view when considering evidence in IDs and can complement each other in order to perform a better decision analysis and to explain the reasoning.[3]

3.3 What-if reasoning: analysis of non-optimal strategies

In Elvira it is possible to have a strategy in which some of the policies are imposed by the user and the others are computed by maximization. The way of imposing a policy consists in setting a probability distribution P_D for the corresponding decision D by means of Elvira's

[3]In the future, we will implement in Elvira Ezawa's method and the possibility of computing the expected value of perfect information (EVPI).

GUI; the process is identical to editing the conditional probability table of a chance node. In fact, such a decision will be treated by Elvira as it were a chance node, and the maximization is performed only on the rest of the decision nodes.

This way, in addition to computing the optimal strategy (when the user has imposed no policy), as any other software tool for influence diagrams, Elvira also permits to analyze how the expected utilities and the rest of the policies would vary if the decision maker chose a non-optimal policy for some of the decisions (what-if reasoning).

The reason for implementing this explanation facility is that when we were building a certain medical influence diagram (Luque et al., 2005) our expert wondered why the model recommended not to perform a certain test. We wished to compute the a posteriori probability of the disease given a positive result in the test, but we could not introduce this "evidence", because it was incompatible with the optimal policy (not to test). After we implemented the possibility of imposing non-optimal policies (in this case, performing the test) we could see that the posterior probability of the disease remained below the treatment threshold even after a positive result in the test, and given that the result of the test would be irrelevant, it is not worthy to do it.

3.4 Decision trees

Elvira can expand a decision tree corresponding to an ID, with the possibility of expanding and contracting its branches to the desired level of detail. This is especially useful when building IDs in medicine, because physicians are used to thinking in terms of scenarios and most of them are familiar with decision trees, while very few have heard about influence diagrams. One difference of Elvira with respect to other software tools is the possibility of expanding the nodes in the decision tree to explore the decomposition of its utility given by the structure of super-value nodes of the ID.

3.5 Sensitivity analysis

Recently Elvira has been endowed with some well-known sensitivity analysis tools, such as one-way sensitivity analysis, tornado diagrams, and spider diagrams, which can be combined with the above-mentioned methods for the explanation of reasoning. For instance, given the ID in Figure 1, one-way sensitivity analysis on the prevalence of the disease X can be used to determine the threshold treatment, and we can later see whether the result of test Y makes the probability of X cross that threshold. In the construction of more complex IDs this has been useful for understanding why some tests are necessary or not, and why sometimes the result of a test is irrelevant.

4 Conclusions

The explanation of reasoning in expert systems is necessary for debugging the knowledge base, for facilitating their acceptance by human users, and for using them as tutoring systems. This is especially important in the case of influence diagrams, because inference in probabilistic graphical models seems to have little relation with human thinking. Nevertheless, in general current software tools for building and evaluating decision trees and IDs offer very little support for analyzing the "reasoning": usually they only permit to compute the value of information, to perform some kind of sensitivity analysis, or to expand a decision tree, which can be hardly considered as explanation capabilities.

In this paper we have described some facilities implemented in Elvira that have been useful for understanding the knowledge contained in a certain ID, and why its evaluation has led to certain results, i.e., the optimal strategy and the expected utility. They can analyze how the posterior probabilities, policies, and expected utilities would vary if the decision maker applied a different (non-optimal) strategy. Most of these explanation facilities are based on the construction of a so-called Cooper policy network, which is a true Bayesian network, and consequently all the explanation options that were implemented for Bayesian networks in Elvira are also available for IDs, such as the possibility of handling several evidence cases simultaneously. Another novelty of Elvira with respect to most software tools for IDs is the ability to include super-value nodes in the IDs, even when expanding the decision tree equivalent to an ID.

We have also mentioned in this paper how these explanation facilities, which we have developed as a response to the needs that we have encountered when building medical models, have helped us to explain the IDs to the physicians collaborating with us and to debug those models (see (Lacave et al., 2006a) for further details).

Acknowledgments

Part of this work has been financed by the Spanish CICYT under project TIC2001-2973-C05. Manuel Luque has also been partially supported by Department of Education of the Comunidad de Madrid and the European Social Fund (ESF).

References

G. F. Cooper. 1988. A method for using belief networks as influence diagrams. In *Proceedings of the 4th Workshop on Uncertainty in AI*, pages 55–63, University of Minnesota, Minneapolis.

R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter. 1999. *Probabilistic Networks and Expert Systems*. Springer-Verlag, New York.

F. J. Díez, J. Mira, E. Iturralde, and S. Zubillaga. 1997. DIAVAL, a Bayesian expert system for echocardiography. *Artificial Intelligence in Medicine*, 10:59–73.

F.J. Díez. 2004. Teaching probabilistic medical reasoning with the Elvira software. In R. Haux and C. Kulikowski, editors, *IMIA Yearbook of Medical Informatics*, pages 175–180. Schattauer, Sttutgart.

The Elvira Consortium. 2002. 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.

K. Ezawa. 1998. Evidence propagation and value of evidence on influence diagrams. *Operations Research*, 46:73–83.

F. V. Jensen. 2001. *Bayesian Networks and Decision Graphs*. Springer-Verlag, New York.

C. Lacave and F. J. Díez. 2002. A review of explanation methods for Bayesian networks. *Knowledge Engineering Review*, 17:107–127.

C. Lacave, R. Atienza, and F. J. Díez. 2000. Graphical explanation in Bayesian networks. In *Proceedings of the International Symposium on Medical Data Analysis (ISMDA-2000)*, pages 122–129, Frankfurt, Germany. Springer-Verlag, Heidelberg.

C. Lacave, M. Luque, and F. J. Díez. 2006a. Explanation of Bayesian networks and influence diagrams in Elvira. Technical Report, CISIAD, UNED, Madrid.

C. Lacave, A. Oniśko, and F. J. Díez. 2006b. Use of Elvira's explanation facility for debugging probabilistic expert systems. *Knowledge-Based Systems*. In press.

M. Luque, F. J. Díez, and C. Disdier. 2005. Influence diagrams for medical decision problems: Some limitations and proposed solutions. In J. H. Holmes and N. Peek, editors, *Proceedings of the Intelligent Data Analysis in Medicine and Pharmacology*, pages 85–86.

D. Nilsson and F. V. Jensen. 1998. Probabilities of future decisions. In *Proceedings from the International Conference on Informational Processing and Management of Uncertainty in knowledge-based Systems*, pages 1454–1461.

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

J. W. Wallis and E. H. Shortliffe. 1984. Customized explanations using causal knowledge. In B. G. Buchanan and E. H. Shortliffe, editors, *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, chapter 20, pages 371–388. Addison-Wesley, Reading, MA.

M. P. Wellman. 1990. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, 44:257–303.