SEQUENTIAL DECISION PROBLEMS REPRESENTED BY SET-VALUED INFLUENCE DIAGRAMS

Daniel Kikuti*, Fabio Gagliardi Cozman*

**Av. Professor Mello Moraes, 2231*
Escola Politécnica, Universidade de São Paulo
São Paulo, SP, Brazil

Emails: daniel.kikuti@poli.usp.br, fgcozman@usp.br

Abstract— This paper presents an algorithm that evaluates an extended version of influence diagrams. Chance variables represent uncertainty about states of the nature through a set of probability measures. This allows us to model prior ignorance in probability inferences, by not forcing the expert to express knowledge through a single and precise distribution. Such a representation also implies that there may exist many expected values for each possible action. In this context, the decision maker must use some criterion to compare the actions in such set. In this paper we discuss two criteria (Γ-maximix and E-admissibility) through an example.

Keywords— Influence diagrams, decision making, uncertainty, credal sets, criteria of choice.

Resumo— Este artigo apresenta um algoritmo que resolve uma versão estendida de diagramas de influência. Variáveis de chance representam incertezas dos estados da natureza através de um conjunto de medidas de probabilidade. Isto nos permite modelar ignorância a priori em inferências probabilísticas, não forçando o especialista a expressar seu conhecimento através de uma única e precisa distribuição. Esta representação também implica que podem existir diversos valores esperados para cada ação. Neste contexto, o tomador de decisão deve utilizar algum critério para comparar as ações em tal conjunto. Neste artigo detalhamos dois critérios (Γ-maximix and E-admissibility) através de um exemplo.

Palavras-chave— Diagramas de influência, tomada de decisão, incerteza, conjuntos credais, critérios de escolha.

1 Introduction

A obvious characteristic of intelligent systems is the ability to make autonomous decisions even in uncertain settings. Decision theory provides a formal framework for determining optimal actions; the notion of optimal is allowed to have a number of different meanings, the most common being the maximization of the agent's expected utility. However, quite often uncertain scenarios are too complex to be adequately described by a precise probability distribution — imprecise probabilities may arise from an incomplete understanding of a decision situation, lack of prior knowledge or empirical data, disagreements between experts, or lack of resources for a complete elicitation procedure (Walley, 1991). In these cases, automatic use of maximization of expected utility may lead to deceptive conclusions (Ellsberg, 1961; Schervish et al., 2003; Seidenfeld, 2004).

In this paper we are interested in using *influence diagrams* to model sequential decision problems where uncertainty is not necessarily represented by a single probability distribution. There are several algorithms in the literature to evaluate influence diagrams in an efficient manner. To *evaluate* or to *solve* an influence diagram is to find a sequence (or sequences) of decisions that satisfies some optimality criterion. The naive solution for solving an influence diagram (by transforming it into a decision tree) is unattractive. More sophisticated algorithms directly evaluate influence diagrams (Shachter, 1986; Shenoy, 1992; Jensen et al., 1994) or reduce it to *Bayesian network* inference problems (Cooper, 1988; Shachter, 1998; Nielsen, 2001; Shachter, 1999). In this paper we focus on the algorithm presented in (Nielsen, 2001; Shachter, 1999), extending it to deal with uncertainty represented by set of probabilities.

In the presence of sets of probability distributions, one may think of maximizing either the minimum expected utility (Γ-*maximin*) (Berger, 1985; Gilboa and Schmeidler, 1989), or the maximum expected utility (Γ-*maximax*) (Satia and Jr., 1973), or a mixture of both according to a "caution variable" representing the degree of ambiguity aversion (Γ-*maximix*) (Utkin and Augustin, 2005). Another possibility is to consider the set of admissible actions, from which we eliminate the dominated (non-optimal) ones. Accordingly, some criteria of admissibility are: *interval dominance*, *maximality* and *E-admissibility*. Algorithms for these criteria can be found in (Kikuti et al., 2005; Utkin and Augustin, 2005).

Our main contribution in this paper is an algorithm to evaluate influence diagrams with imprecise probabilities. Section 2 briefly reviews the basics of set of probabilities and influence diagrams; in particular, we are going to detail the relationship between influence diagrams and Credal networks. Section 3 presents the algorithm for evaluating influence diagrams and how it is linked with the criteria of choice. From the criteria above, we present two algorithms: Γ-Maximix and E-Admissibility. An example and the conclusions are presented at Sections 4 and 5 respectively.

2 Background and problem statement

Decision theory with imprecise probabilities is caracterized in the following way: an individual is faced with several alternative courses of *actions* $\mathbb{A} = \{a_1, \ldots, a_m\}$. The consequences of every action depend on the *state of nature* $\Omega = \{\omega_1, \ldots, \omega_n\}$. An outcome c_{ij} is associated for choosing action a_i in state ω_j.

The *uncertainty* (or *ambiguity*[1]) over the states of the nature is represented by *credal sets* (Levi, 1980). A credal set $K(X)$ is a set of probability distributions for a random variable $X : \Omega \to \mathbb{R}$, which takes different possible values on Ω. We assume that all variables are categorical and the credal sets are closed and convex with finitely many vertices — $K(X)$ is represented by a polytope in $\mathbb{R}^n$ space, where n is the cardinality of the event space for variable X. Given a credal set and a real-valued utility function $u : (\mathbb{A} \times \Omega) \to \mathbb{R}$ defined over outcomes, one may compute the set of expectations by:

$$E[u(a_i)] = \sum_{j=1}^{n} p(x_j) u(c_{ij}) \quad (1)$$

Equation 1 results in a set of expected utility values because $p(x_j) \in K(X = x_j)$ may have more than one probability measure associated with it. The notation $p(x_j)$ is a shorthand for $p(\{\omega \in \Omega : X(\omega) = x_j\})$. This set of expectations is usually represented in terms of interval-values $[\underline{E}[u(a_i)], \overline{E}[u(a_i)]]$, where $\underline{E}[u(a_i)]$ is the *lower expectation* $(\min E[u(a_i)])$ and $\overline{E}[u(a_i)]$ is the *upper expectation* $(\max E[u(a_i)])$. Lower and upper probabilities are defined similarly (Giron and Rios, 1980; Walley, 1991).

These concepts describe a single stage decision problem. In this paper we are considering dynamic choice situations, i.e., sequential decision problems represented by influence diagrams where decisions are made after the resolution of some uncertainty.

An influence diagram (Howard and Matheson, 1984) is a graphical representation of uncertainty quantities and decisions that explicitly reveals probabilistic dependence and the flow of information. Previous researches (Shenoy, 1992; Jensen et al., 1994) have applied methods for probabilistic inference into sequential decision problems. In fact, an influence diagram can be seen as a Bayesian network (Pearl, 1988) augmented with decision nodes and value nodes (Cooper, 1988). When we have chance nodes representing credal sets, we can say that the influence diagram is a *credal network* augmented with decision and value nodes. A credal network is a directed acyclic graph where each node of the graph is associated with a variable X_i and with a collection of conditional credal sets $K(X_i|Pa(X_i))$, where $pa(X_i)$ denotes the parents of X_i. A conditional credal set is obtained by applying Bayes rule to every distribution in a credal set. We adopt the following definition of independence, usually referred to as *strong independence*: two variables X and Y are strongly independent when the credal set $K(X, Y)$ has all vertices satisfying stochastic independence of X and Y (that is, all vertices factorize as $P(X)\,P(Y)$) (Couso et al., 2000; Cozman, 2000). A *marginal inference* in a credal network is the computation of lower/upper probabilities in an extension of the network.

An influence diagram is represented by a directed acyclic graph $\mathcal{G} = \{\mathbb{V}, \mathbb{E}\}$ with $\mathbb{V}$ patitioned into three subsets: decision nodes (rectangles), chance nodes (circles or ovals) and value nodes (diamond shapes). Each decision node contains a set of possible actions, each chance node is associated with a set of conditional probability tables and each value node is associated with a utility function. We assume that the influence diagram can have more than one value node in order to allow dynamic programming and that the total utility is the sum of the local utilities (Tatman and Shachter, 1990). Arcs into chance nodes indicate probabilistic dependence (conditional arcs) and arcs into decisions specify the information available at the time of the decision (informational arcs).

In influence diagrams, there must be a total ordering of the decision nodes indicating the order in which the decisions are made $\mathbb{V}_D = \{D_1, \ldots, D_n\}$. As a consequence of this ordering and the set of informational arcs, it is possible to partition the set of chance variables into a collection of disjoint subsets $\mathbb{V}_C = \{\mathbb{C}_0, \ldots, \mathbb{C}_n\}$, where $\mathbb{C}_i$ denotes the chance variables observed between decision D_i and D_{i+1}. This induces the *partial order* $\prec$ on $\mathbb{V}_C \cup \mathbb{V}_D$: $\mathbb{C}_0 \prec D_1 \prec \mathbb{C}_1 \prec \ldots \prec D_n \prec \mathbb{C}_n$. A *policy* δ_i is then defined as a mapping from the past of D_i $(pa(D_i) \cup_{j=0}^{i-2} \mathbb{C}_j)$ to the state space of $dom(D_i)$, i.e., a policy specify an action given all observations made prior to making decision in D_i. A *strategy* s is an ordered set of decision policies $s = \{\delta_1, \ldots, \delta_n\}$ for each $D_i \in \mathbb{V}_D$. The *optimal strategy* $s^* = \{\delta_1^*, \ldots, \delta_n^*\}$ depends on the criteria of optimality adopted by the decision maker. To solve an influence diagram amounts to determine the optimal strategy and compute the expected utility for adhering to this strategy.

The decomposition method of (Nielsen, 2001) proposes the solution to the problem by decomposing the influence diagram into a collection of smaller influence diagrams. The decomposition produces an influence diagram for each decision variable of interest, and each of this influence diagrams contains exactly the variables necessary

[1]Ellsberg defines ambiguity as the vague or unsure probability judgments that a decision maker assigns to a particular problem (Ellsberg, 1961).

and sufficient for determining an optimal policy for the associated decision variable; hence, the influence diagram can be solved independently of each other.

Starting with the last decision, we can obtain the expected utility by solving the Equation 2.

$$E[D_n = d] = \sum_{X_{Q_n}} P(X_{Q_n}|X_{E_n})u(d, X_{Q_n}) \quad (2)$$

The X_{Q_n} in Equation 2 is the set of queried variables, X_{E_n} is the set of evidences (observed variables), and $u(d, X_{Q_n})$ is the utility function. To calculate the expected utility for an optimal strategy strategy s^* we expand the expression above to:

$$E[s^*] = \sum_{\mathbb{C}_0} \underset{\mathrm{D}_1}{opt} \ldots \sum_{\mathbb{C}_{n-1}} \underset{\mathrm{D}_n}{opt} \sum_{\mathbb{C}_n} \prod_{X \in \mathbb{V}_C} P(X|pa(X)) \sum_{U \in \mathbb{V}_U} u(pa(U)) \quad (3)$$

In our algorithms, the strategies are represented by multilinear programs — non linear programs whose objective function and constraints involve the variables through sum of products — and are solved using a method of Reformulation-Linearization (Sherali and Tuncbilek, 1992).

3 An influence diagram algorithm and two criteria of choice

This section presents an algorithm for evaluating influence diagrams based on Shachter's algorithm. This algorithm determines the variables required to solve the influence diagram with relation to the decision node D_i ($\forall i, 1 \leq i \leq n$) (Shachter, 1998; Nielsen, 2001). The Algorithm 1 allows not only a policy for each decision node (as in the case with precise probabilities), but also a set of policies classified as admissible according to a previously defined criterion of optimality. As the result of the evaluation, the algorithm returns a strategy (or a set of strategies).

Lines 1 to 5 decompose the influence diagram $\mathcal{I}$ into a set of smaller influence diagrams. The array $\mathbb{I}$ contains the information necessary and sufficient to evaluate each decision node separately. The function `GetValueNodes` returns the set of value nodes relevant to decision D_i (the set of value nodes to which there exists a directed path, excluding informational arcs, from D_i in $\mathcal{I}$). A proof that this function returns the relevant set can be found in (Nielsen and Jensen, 1999). The same argument can be used to show that this is true also to chance variables with imprecise probabilities. The function `GetD-Connected` returns the set of variables that are d-connected to a variable in V_i given D_i and its predecessors. In this step we prune the unnecessary nodes to evaluate the expectations of the actions in D_i. A discussion of irrelevance and independence relations in credal networks can be found in (Cozman, 2000). The variable Req_i contains the set of variables required for D_i.

Algorithm 1: Evaluate

Input: An influence diagram $\mathcal{I}$
Output: A set of admissible strategies $\mathbb{S}$

1 **foreach** D_i, $i \leftarrow N$ **to** *1* **do**
2   $V_i \leftarrow$ `GetValueNodes`($\mathcal{I}$, D_i);
3   $\mathcal{D}_i \leftarrow$ `GetD-Connected`($\mathcal{I}$, D_i, V_i);
4   $Req_i \leftarrow \mathcal{D}_i \cap$ `Pred`($\mathcal{I}$,D_i);
5   $\mathbb{I}[i] \leftarrow V_i$, $\mathcal{D}_i$ and Req_i ;
6 **foreach** $D_i \in \mathbb{I}$, $i \leftarrow N$ **to** *1* **do**
7   $\Delta \leftarrow \emptyset$;
8   **foreach** *action* $d \in D_i$ **do**
9   **foreach** *configuration* j *of* Req_i **do**
10       $\delta_{i_j} \leftarrow$ Multilinear program with objective function given by Equation 3 s.t. constraints on probability values of variables in $\mathbb{I}[i]$;
11       $\Delta \leftarrow \Delta \cup \delta_{i_j}$;
12   Admissible[i] $\leftarrow$ `Criterion--X`(Δ);
13 $\mathbb{S} \leftarrow$ Enumeration of all admissible strategies;
14 **return** $\mathbb{S}$;

Lines 6 to 12 determine the admissible policies for each smaller influence diagram in $\mathbb{I}$ and evaluate it. The Admissible array at line 12 contains the admissible set of policies determined by the choosen criterion. The function Criterion—X is a generic function that must be replaced by some properly implemented criteria of optimality. If we choose an ordering criterion, we always have a policy for each decision node; if we choose an admissibility criterion, we would have in the worst case $m * n$ policies to consider, where m is the number of possible actions available at the decision node and n is the number of all possible configurations of the necessary variables to the decision node. At line 13 we have the enumeration of all strategies $s = \{\delta_1, \ldots, \delta_n\}$, i.e., we take one policy for each decision node and builds a strategy.

The subsections below give more details about two of criteria of optimality. Both criteria demand to solve a linear number of multilinear programs, which computational complexity depends on the structure of problem.

3.1 Γ-Maximix criterion

With imprecision on probability values, we have an interval of expectations such that we can adopt a criterion of choice that looks only at the lower expectation (Γ-*Maximin*), or looks only at the upper expectation (Γ-*Maximax*) or looks at both, lower and upper expectations, suggesting a decision based on a mixture of both values according to a *caution parameter* η reflecting the degree of ambiguity aversion (Γ-*Maximix*) (Utkin and Augustin, 2005).

The optimal choice δ^* for Γ-Maximix is:

$$\delta^* = \max(\eta \underline{E}[\delta_i] + (1-\eta)\overline{E}[\delta_i])$$

The more ambiguity averse the decision maker is, the higher is the influence of the lower interval limit of the generalized expected utility[2] (Utkin and Augustin, 2005).

This criterion is described in algorithm below.

Algorithm 2: Criterion–Γ-Maximix

Input: A set of policies Δ and a parameter η reflecting the degree of ambiguity aversion
Output: The optimal strategy δ^*

$\delta^* \leftarrow null$;
hme $\leftarrow -\infty$;
foreach $\delta_i \in \Delta$ **do**
  mixed $\leftarrow \eta \underline{E}[\delta_i] + (1-\eta)\overline{E}[\delta_i]$;
  if mixed > hme **then**
    $\delta^* \leftarrow \delta_i$;
    hme $\leftarrow$ mixed;
return δ^*;

3.2 *E-admissibility criterion*

Now let's consider that the set of actions cannot be completely ordered. Then it is possible to say that one action is admissible if its upper expectation is greater than or equal to the greatest value of the lower expectation among the actions (this criterion is known as *interval dominance*); or we can say that one is admissible in a pairwise comparison if it is not worst for all probability value (*maximality*); or we can adopt the concept of *E-admissibility*, which restricts the decision maker's admissible choices to those that are Bayes for at least one probability measure P in the relevant credal sets. That is, given a choice set Δ of feasible policies and a credal set K representing imprecise beliefs, the policy $\delta \in \Delta$ is E-admissible when, for at least one $P \in K$, δ maximizes expected utility (Schervish et al., 2003):

For each policy δ, we are interested in finding a probability distribution for which δ is optimal in the standard expected utility sense. If this probability distribution exists, then δ is E-admissible. In other words, the policy $\delta_i \in \Delta$ is E-admissible if there exists a $P \in K$ such that for all $\delta_j \in \Delta$, $\delta_j \neq \delta_i$, we have $E[\delta_i - \delta_j] \geq 0$. These (multilinear) constraints must all be satisfied to show that δ_i is E-admissible; if the constraints cannot be satisfied, then δ_i is not E-admissible. We thus obtain the following algorithm, where LR is a list of constraints produced by pairs of strategies:

Algorithm 3: Criterion–E-Admissibility

Input: A set of policies Δ
Output: A set of admissible policies

1 $n \leftarrow$ Number of policies in Δ ;
2 **for** $i \leftarrow 1$ **to** n **do**
3   LR $\leftarrow$ null;
4   **for** $j \leftarrow 1$ **to** n **do**
5   **if** $i \neq j$ **then**
6   LR $\leftarrow$ LR $\cup\, E[\delta_i - \delta_j] \geq 0$;
7   $Q \leftarrow$ set of all constraints on probability values plus LR;
8   $P \leftarrow \arg\max_p E[\delta_i]$ s.t. constraints on Q;
9   **if** $P \neq null$ **then**
10     $\delta_i.admissible \leftarrow$ true;
11   **else**
12     $\delta_i.admissible \leftarrow$ false;
13 **return** All alternatives not marked as false;

Lines 3 to 6 generate all constraints that are required to satisfy E-admissibility of a policy δ_i, and line 7 collects constraints on probabilities. Line 8 requires the solution of a multilinear program. The whole algorithm depends on the number of policies, and not directly on the number of distributions in the credal sets. Even though the properties of the credal sets certainly affect the solution of the relevant multilinear programs, there is no need to represent the credal sets explicitly, or to enumerate their vertices (necessary steps in algorithms previous than (Kikuti et al., 2005; Utkin and Augustin, 2005). In a sense, the complexity of credal sets is "hidden" within the multilinear programs.

4 Example

Consider the classic *oil wildcatter* problem (Raiffa, 1968) depicted in Figure 1. An oil wildcatter must decide whether to drill or not to drill. The cost of drilling is $\$70K$. If he decides to drill, the hole may be soaking, wet or dry (with a return of $\$270K$, $\$120K$ or $\$0$ respectively). Suppose that the prior probabilities for soaking, wet and dry are given as interval-values: [0.181, 0.222], [0.333, 0.363] and [0.444, 0.454] respectively. At the cost of $\$10K$, the oil wildcatter could decide to take a seismic soundings of the geological structure at the site. The specifics of the test are given in Table 1.

Applying Algorithm 1, we have the influence

[2]Note that $\eta = 1$ corresponds to strict ambiguity aversion (Γ-Maximin) and a $\eta = 0$ corresponds to a maximal ambiguity seeking (Γ-Maximax)

Test Results (S)	Amount of Oil (O)		
	dry	wet	soaking
no (ns)	0.6	0.3	0.1
open (os)	0.3	0.4	0.4
closed (cs)	0.1	0.3	0.5

Table 1: Probabilities of seismic test results conditional on amount of oil.

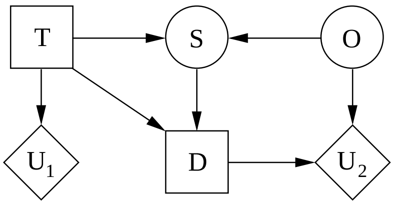

Figure 1: Influence diagram for the oil wildcatter problem.

diagrams depicted in Figure 2 to evaluate (information at array $\mathbb{I}$). The decision `T` represents the decision of taking or not the seismic test and it is to be made based on no information, that is, we have just two policies to analyse ($\delta_T = yes$ and $\delta_T = no$). The decision `D` represents the decision of drilling and it is to be made conditional on `S` and on `T` (required nodes are shaded), thus a policy for `D` specify how to act in 6 possible configurations of the required nodes.

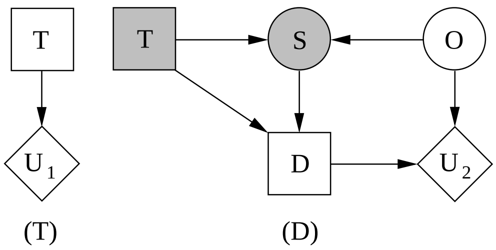

Figure 2: Decomposition of the influence diagram in Figure1 into two smaller influence diagrams.

The expectation to δ_D is obtained by the multilinear program below:

$$EU[\delta_D] = \sum_S P(S) \underset{D}{opt} \sum_O \frac{P(O)P(S|O,T)}{P(S)} U(D,O) \text{ s.t.: constraints on P(O) and P(S|O,T)}$$

Table 2 shows the expected utility function over variables `T`, `S` and `D`.

Applying the criterion Γ-Maximix with $\eta = 1.0$ (we are just looking the lower expectations), we have that the optimal policy is to take the seismic test and not drill only if it indicates no structure. The local expected utility computation for

D	S	T	lower	upper
no	–	–	0	0
yes	–	no	21.77	30.77
yes	ns	yes	-24.85	-24.01
yes	os	yes	27.22	30.05
yes	cs	yes	81.04	86.85

Table 2: Lower and upper expected utility for `D`

δ_D^* is $\approx \$28.87K$. The total utility is the sum of the local utilities. For the optimal strategy $s^* = \{\delta_T^*, \delta_D^*\}$ we have $EU[s^*] = EU(\delta_T^*) + EU[\delta_D^*] = -10K + 28.87K = \$18.87K$.

Applying the criterion E-admissibility we have obtained the same optimal strategy (coincidentally). This can be easily verified if we note that the constraints of the policies suggesting not to drill cannot maximize the expected utility for any probability value (except for no structure in seismic test). The expectations for this criterion is $\underline{E}[s^*] = \$18.86K$ and $\overline{E}[s^*] = \$24.16K$.

5 Conclusion

In this paper we have presented an algorithm to solve influence diagrams where chance nodes are associated with sets of probabilities. The algorithm is based on graph-theoretic properties of Bayesian and credal networks. The algorithm solves the original influence diagram by decomposing it into smaller and independent influence diagrams.

The algorithm was specialized to deal with optimality criteria: Γ-maximix and E-admissibility. Both criteria require use of multilinear programming technique for finding optimal actions. In the E-admissibility criteria the decision maker is allowed to have more than one optimal action; this may be computational costly if the number of admissible actions in each decision node is large and the influence diagram is decomposed in many sub diagrams.

Acknowledgments

This work has been financially supported by FAPESP grant 2004/09568-0; the first author is supported by FAPESP grant 03/11165-9 the second author is partially supported by CNPq grant 3000183/98-4.

References

Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*, Springer-Verlag, New York.

Cooper, G. F. (1988). A method for using belief networks as influence diagrams, *Proceedings* *of the Fourth Conference on Uncertainty in Artificial Intelligence*, Minneapolis, pp. 55–63.

Couso, I., Moral, S. and Walley, P. (2000). A survey of concepts of independence for imprecise probabilities, *Risk, Decision and Policy* **5**: 165–181.

Cozman, F. G. (2000). Separation properties of sets of probabilities, *in* C. Boutilier and M. Goldszmidt (eds), *XVI Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, San Francisco, pp. 107–115.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms, *Quarterly Journal of Economics* **75**(4): 643–669.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics* **18**(2): 141–153.

Giron, F. J. and Rios, S. (1980). *Quasi-Bayesian behaviour: a more realistic approach to decision making?*, University Press, Valencia.

Howard, R. A. and Matheson, J. E. (1984). Influence diagrams, *The Principles and Applications of Decision Analysis*, editors, Menlo Park, CA, pp. 719–762.

Jensen, F., Jensen, F. V. and Dittmer, S. L. (1994). From influence diagrams to junction trees, *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence*, Seattle, Washington, pp. 367–373.

Kikuti, D., Cozman, F. G. and de Campos, C. P. (2005). Partially ordered preferences in decision trees: computing strategies with imprecision in probabilities, *IJCAI Workshop about Advances on Preference Handling*, Edinburgh, United Kingdom, pp. 118–123.

Levi, I. (1980). *The enterprise of Knowledge*, MIT Press, Massachusetts.

Nielsen, T. D. (2001). Decomposition of influence diagrams, *ECSQARU '01: Proceedings of the 6th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, Springer-Verlag, London, UK, pp. 144–155.

Nielsen, T. D. and Jensen, F. V. (1999). Welldefined decision scenarios, *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, Stockholm, Sweden, pp. 502–511.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann Publishers, Los Altos, CA.

Raiffa, H. (1968). *Decision analysis: introductory lectures on choices under uncertainty*, Addison-Welsey, Massachusetts.

Satia, J. K. and Jr., R. E. L. (1973). Markovian decision processes with uncertain transition probabilities, *Operations Research* **21**(3): 728–740.

Schervish, M. J., Seidenfeld, T., Kadane, J. B. and Levi, I. (2003). Extensions of expected utility theory and some limitations of pairwise comparisons, *Proceedings of 3rd International Symposium on Imprecise Probabilities and Their Applications*, Lugano, Switzerland, pp. 496–510.

Seidenfeld, T. (2004). A contrast between two decision rules for use with (convex) sets of probabilities: Gamma-maximin versus E-admissibility, *Syntese* **140**(3): 69–88.

Shachter, R. (1998). Bayes-ball: the rational pasttime (for determining irrelevance and requisite information in belief networks and influence diagrams, *Proceedings of the 14th Annual Conference on Uncertainty in Artificial Intelligence (UAI-98)*, Morgan Kaufmann, San Francisco, CA, pp. 480–48.

Shachter, R. (1999). Efficient value of information computation, *Proceedings of the 15th Annual Conference on Uncertainty in Artificial Intelligence (UAI-99)*, Morgan Kaufmann, San Francisco, CA, pp. 594–60.

Shachter, R. D. (1986). Evaluating influence diagrams, *Operations Research* **34**(6): 871–882.

Shenoy, P. P. (1992). Valuation-based systems for bayesian decision analysis, *Operations Research* **40**(3): 463–484.

Sherali, H. D. and Tuncbilek, C. H. (1992). A global optimization algorithm for polynomial programming problems using a reformulation-linearization technique, *Journal of Global Optimization (Historical Archive)* **2**(1): 101–112.

Tatman, J. A. and Shachter, R. D. (1990). Dynamic programming and influence diagrams, *IEEE Transactions on systems, man and cybernetics* **20**(2): 365–379.

Utkin, L. V. and Augustin, T. (2005). Powerful algorithms for decision making under partial prior information and general ambiguity attitudes, *Proceedings of 4th International Symposium on Imprecise Probabilities and Their Applications*, Pittsburgh, Pennsylvania.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities*, Chapman and Hall, Londom.