

Variable Elimination for Interval-Valued Influence Diagrams

Rafael Cabañas¹, Alessandro Antonucci²,
Andrés Cano¹, and Manuel Gómez-Olmedo¹

¹ Department of Computer Science and Artificial Intelligence
CITIC, University of Granada, Spain
{rcabanas, acu, mgomez}@decsai.ugr.es

² Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA)
Galleria 2, Manno-Lugano, Switzerland
alessandro@idsia.ch

Abstract. Influence diagrams are probabilistic graphical models used to represent and solve decision problems under uncertainty. Sharp numerical values are required to quantify probabilities and utilities. Yet, real models are based on data streams provided by partially reliable sensors or experts. We propose an interval-valued quantification of these parameters to gain realism in the modelling and to analyse the sensitivity of the inferences with respect to perturbations of the sharp values. An extension of the classical influence diagrams formalism to support interval-valued potentials is provided. Moreover, a variable elimination algorithm especially designed for these models is developed and evaluated in terms of complexity and empirical performances.

Keywords: Influence diagrams, Bayesian networks, credal networks, sequential decision making, imprecise probability.

1 Introduction

Influence diagrams are probabilistic graphical models able to cope with decision problems with uncertainty. The parameters of an influence diagram are conditional probabilities for single variables given some other variables, or utilities depending on given sets of variables. The quantification of these parameters is based on a statistical processing of data or on the elicitation of expert knowledge.

Exactly as Bayesian networks, influence diagrams require sharp estimates of their parameters. Yet, when coping with expert knowledge, sharp values can be unfit to express judgements (e.g., which is the number modelling the probability for an option *more probable* than its negation?). This issue appears also when coping with scarce or missing data (e.g., probabilities conditional on rare events).

For reasons of this kind, in the last two decades, various extensions of Bayesian networks to support more general probabilistic statements have been proposed. These models have been developed in the field of possibility theory [2], evidence theory [17], and imprecise probability [4]. The latter models, called *credal networks*, offer a direct sensitivity analysis interpretation: a credal network is a collection of Bayesian networks, all over the same variables and with the same graph, whose parameters are consistent with constraints (e.g., intervals) modelling a limited ability in the assessment of sharp estimates. Similarly, various

extensions of decision trees have also been proposed [9, 11]. The situation is different for influence diagrams. The early attempts of Fertig and Breese [7] first, and Zaffalon [6] after, to extend these models to non-sharp quantification are to the best of our knowledge the only works in this direction.³ This is unfortunate as the above considerations about the difficulty of assessing sharp estimates for probabilities are even more compelling for utilities, which are supposed to model intrinsically qualitative objects such as preferences.

In this paper we extend to the interval-valued case the formalism of influence diagrams by keeping the same sensitivity-analysis interpretation of credal networks: a generalized influence diagram is a collection of classical influence diagrams consistent with the interval constraints. When coping with interval-valued utilities we might have overlaps between the different expectations. In these cases we adopt a conservative approach which rejects decisions leading to certainly dominated options and keeps all the other ones. A first example of *variable elimination* to compute inferences in these generalized models (arc reversal was considered in [7]) is also proposed together with some preliminary tests.

2 Basics

Let us first define the basic notation. We use upper-case letters for variables and lower-case for states. Given a variable X , x is an element of the domain of X , which we denote as Ω_X . Given a set of n variables $\mathbf{X} := (X_1, \dots, X_n)$, and a multi-valued index $J \subseteq [1, n]$, X_J is the joint variable including any X_i such that $i \in J$. Thus, $\Omega_{X_J} = \times_{i \in J} \Omega_{X_i}$, where \times is the Cartesian product. Given a second index I , notation $x_I \sim x_J$ is used to express *consistency*, i.e., to denote the fact that the two states have the same values on $X_{I \cap J}$. *Chance* variables are those whose actual value might be unknown, *decision* variables are those whose actual value can be set by the decision maker. A *potential* over X_J is a map $\psi : \Omega_{X_J} \rightarrow \mathbb{R}$. *Probability potentials* (PP, also called conditional probability tables) are special potentials. Given two disjoint set of variables X_I and X_J , $\phi(X_I, X_J)$ is a PP over X_I given X_J if and only if it is a nonnegative potential such that $\sum_{x_I \in \Omega_{X_I}} \phi(x_I, x_J) = 1$ for each $x_J \in \Omega_{X_J}$.

Influence Diagrams. Influence diagrams (IDs) [8] are a class of graphical models designed to formalize sequential decision problems with uncertainty. The uncertainty is represented by PPs, while the user preferences are represented by generic potentials called here *utility potentials* (UPs).

An ID over a set of variables (\mathbf{X}, \mathbf{D}) is made of a qualitative and a quantitative part. The qualitative part is a directed acyclic graph (DAG) \mathcal{G} with three types of nodes. *Chance* nodes are depicted as circles and are in one-to-one correspondence with the chance variables, i.e., the variables in \mathbf{X} . Decision nodes are depicted as squares and associated to decision variables \mathbf{D} . Utility nodes are depicted as a diamonds and should be barren. For chance and decision variables

³ Sensitivity analysis does not require the specification of more general class of models, being only focused on the results of the inferences. Thus, it should be regarded as a different topic, which, as a matter of fact, received more attention (e.g., [14]).

we use the terms node and variable interchangeably. Utility nodes are not associated to variables. Still, these nodes are jointly denoted as U . The immediate predecessors of a node Y according to the \mathcal{G} are called *parents* and denoted as Π_Y . From the quantitative point of view, for each chance node, a PP over the corresponding variable and its parents is defined, while, for each utility node, a UP potential over the parents should be assessed. The formal definition of ID is the following.

Definition 1. An influence diagram is a tuple $\langle \mathcal{G}, \mathbf{X}, \mathbf{D}, \mathbf{U}, \Phi, \Psi \rangle$, where \mathcal{G} is a DAG over $\mathbf{X} \cup \mathbf{D} \cup \mathbf{U}$, while $\Phi = \{\phi(X, \Pi_X)\}_{X \in \mathbf{X}}$ and $\Psi = \{\psi(\Pi_U)\}_{U \in \mathbf{U}}$ are collections of, respectively, PPs and UPs.

To model sequential decision problems with IDs, some additional information should be provided. A complete order of decision variables (e.g., $D_1 \prec \dots \prec D_n$), and a partial order for $\mathbf{X} \cup \mathbf{D}$ consistent with the partial one for \mathbf{D} should be formulated. Accordingly the partial order has form $\mathcal{I}_0 \prec D_1 \prec \mathcal{I}_1 \prec \dots \prec D_n \prec \mathcal{I}_n$, with $\cup_{j=0}^n \mathcal{I}_j = \mathbf{X}$ and $\mathcal{I}_i \cap \mathcal{I}_j = \emptyset$ for each i, j . This reflects a temporal interpretation: the chance variables in \mathcal{I}_i are observed before decision D_{i+1} is taken, and the ordering over \mathbf{D} reflects the order in which the different decisions are taken. An ID of this kind is called *regular*. *Non-forgetting assumption* is usually required as well: previous decisions and observations are known at each decision. Here we only consider regular IDs with the non-forgetting assumption. A classical example is here below.

Example 1 (The oil wildcatter [15, 16]). An oil wildcatter must decide whether to drill or not. He is uncertain whether the amount of oil (O) in the place is empty (e), wet (w) or soaking (s). The wildcatter can make seismic tests (S) that will give a closed reflection pattern (c) indicating much oil, an open pattern (o) indicating for some oil, or a diffuse pattern (d) denoting almost no hope for oil. These two are chance variables, while the decision variables are T [to test (t) or not (nt)] and D [to drill (d) or not (nd)]. The utility nodes P and C describe the profit possibly obtained from the presence of oil and the cost of the tests. The DAG of an ID modelling this problem is in Fig. 1. Decision T precedes decision D , while the partial order is complete being $T \prec \{S\} \prec D \prec \{O\}$.

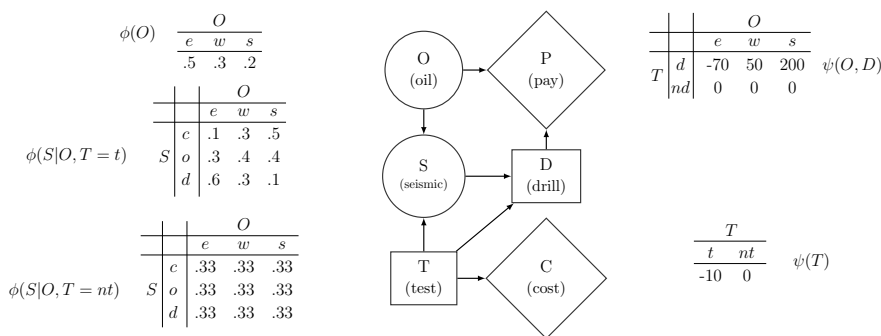


Fig. 1. The oil wildcatter's decision problem as an ID

ID Evaluation. A *policy* for a decision variable D is a map $\delta_D : \Omega_{\Pi_D} \rightarrow \Omega_D$, i.e., a rule to assign D on the basis of the values of the parents. A *strategy* Δ is a collection of policies, one for each $D \in \mathbf{D}$. From a joint observation \mathbf{x} of the chance variables and a strategy Δ , we deduce a joint specification of the decision variables, and hence of the values of the utilities, which are assumed additive. Accordingly, we can compute the expected utility of a strategy as follows:

$$EU(\Delta) := \sum_{\mathbf{x}} \prod_{X \in \mathbf{X}} \phi(x, \pi_X) \sum_{U \in \mathbf{U}} \psi(\pi_U), \quad (1)$$

with $x, \pi_X, \pi_U \sim \mathbf{x}$, and the values of π_X and π_U also consistent with the values of \mathbf{D} obtained from \mathbf{x} and Δ . A typical task with IDs is to detect the strategy Δ^* maximizing the expected utility. This should respect the partial order, i.e.,

$$EU(\Delta^*) = \sum_{\mathcal{I}_0} \max_{D_1} \cdots \max_{D_n} \sum_{\mathcal{I}_n} \prod_{X \in \mathbf{X}} \phi(X | \Pi_X) \sum_{U \in \mathbf{U}} \psi(\Pi_U) \quad (2)$$

The components of this optimal strategy, are called *optimal policies*, and they can be regarded as the intermediate steps of the above maximization. The optimal policy $\delta_{D_i}^*$ associated to D_i is therefore:

$$\delta_{D_i}^*(\Pi_{D_i}) = \arg \max_{D_i} \sum_{\mathcal{I}_i} \max_{D_{i+1}} \cdots \max_{D_n} \sum_{\mathcal{I}_n} \prod_{X \in \mathbf{X}} \phi(X | \Pi_X) \sum_{U \in \mathbf{U}} \psi(\Pi_U). \quad (3)$$

Example 2. In the oil wildcutter’s problem (Ex. 1), when seismic tests have been done and returned a closed reflection pattern the right thing to do is to drill, i.e. $\delta_D^*(S = c, T = t) = d$ as $EU_D(S = c, T = t) = 77.5$.

Variable Elimination. *Variable elimination* (VE) is a typical approach to inference in graphical model. VE algorithms for IDs [10, 18] are commonly used to solve Eq. (2). Unlike VE for Bayesian networks, in regular IDs the elimination order is not arbitrary: it should be the inverse of an order consistent with the partial order associated to the ID [12]. Furthermore, while chance variables are removed by sum, decision variables are instead eliminated by maximization. To describe VE, we first show how to remove a single variable Y in Alg. 1. The whole algorithm consists in iterating the procedure over $\mathbf{X} \cup \mathbf{D}$. When the last variable is eliminated, the algorithm returns a potential with no arguments (i.e., a constant) with the value in Eq. (2).

The operator `dom` returns the variables in the argument of a potential. Sums in line 4 and maxima in line 6 of Alg. 1 are two different forms (the first for decision, the second for chance variables) of *marginalization*, removing the variable from the argument of the potential. The division (line 4) is defined element-wise. For the PP in line 6, Y can be eliminated by instantiating an arbitrary value [12]. When eliminating a decision variable, the maximization of the UP also gives the corresponding optimal policy (line 7).

Algorithm 1 Removing a single (chance or decision) variable Y

```

1:  $(\Phi_Y, \Psi_Y) \leftarrow (\{\phi \in \Phi \mid Y \in \text{dom}(\phi)\}, \{\psi \in \Psi \mid Y \in \text{dom}(\psi)\})$  ▷ Select
2:  $(\phi_Y, \psi_Y) \leftarrow (\otimes_{\phi \in \Phi_Y} \phi, \otimes_{\psi \in \Psi_Y} \psi)$  ▷ Combine
3: if  $Y \in \mathbf{X}$  then
4:    $(\phi'_Y, \psi'_Y) \leftarrow (\sum_Y \phi_Y, \frac{\sum_Y \phi_Y \otimes \psi_Y}{\sum_Y \phi_Y})$  ▷ Remove by sum (chance vars)
5: else
6:    $(\phi'_Y, \psi'_Y) \leftarrow (\phi_{Y=y}, \max_Y \psi_Y)$  ▷ Remove by max (decision vars)
7:    $\delta_Y^* \leftarrow \arg \max_Y \psi_Y$  ▷ Optimal policy (as a byproduct)
8: end if
9:  $(\Phi, \Psi) \leftarrow (\Phi \setminus \Phi_Y \cup \{\phi'_Y\}, \Psi \setminus \Psi_Y \cup \{\psi'_Y\})$  ▷ Update
10: return  $(\Phi, \Psi)$ 

```

The operator \otimes in Alg. 1 combines pairs of potentials as explained here below. It is easy to check that these definitions are well-posed and that the operator is associative and commutative.

- i) given two UPs, say $\psi(X_I)$ and $\psi'(X_J)$, their combination $\psi \otimes \psi'$ is a UP over $X_{I \cup J}$ obtained by element-wise sums, i.e., $(\psi \otimes \psi')(x_{I \cup J}) := \psi(x_I) + \psi'(x_J)$ for each $x_{I \cup J} \in \Omega_{X_{I \cup J}}$, with $x_I, x_J \sim x_{I \cup J}$;
- ii) given a PP $\phi(X_I, X_J)$ and a UP $\psi(X_K)$, their combination $\phi \otimes \psi$ is a UP over $X_L := X_{I \cup J \cup K}$ defined by element-wise products, i.e., $(\phi \otimes \psi)(x_{I \cup J \cup K}) := \phi(x_I, x_J) \cdot \psi(x_K)$, for each $x_{I \cup J \cup K} \in \Omega_{X_{I \cup J \cup K}}$, with $x_I, x_J, x_K \sim x_{I \cup J \cup K}$;
- iii) finally, given two PPs, say $\phi(X_I, X_J)$ and $\phi'(X_K, X_L)$, their combination $\phi \otimes \phi'$ is a PP over $X_{I \cup K}$ given $X_{(J \cup L) \setminus (I \cup K)}$ defined by element-wise products, i.e., $(\phi \otimes \phi')(x_{I \cup K}, x_{(J \cup L) \setminus (I \cup K)}) := \phi(x_I, x_J) \cdot \phi'(x_K, x_L)$ for each $x_{I \cup K} \in \Omega_{X_{I \cup K}}$ and $x_{(J \cup L) \setminus (I \cup K)} \in \Omega_{X_{(J \cup L) \setminus (I \cup K)}}$, with $x_I, x_J, x_K, x_L \sim x_{I \cup K}, x_{(J \cup L) \setminus (I \cup K)}$.

3 Interval-Valued Potentials

The main goal of this paper is to extend IDs to support interval-valued specifications. To do that, we first formalize the basic notion of interval-valued potential.

Definition 2. An interval-valued utility potential (IUP) over X_I is a pair of UPs over X_I . We use the compact notation $\underline{\psi}(X_I)$ for a IUP over X_I , $\underline{\psi}$ and $\bar{\psi}$ are the two UPs involved in the specification and are called, respectively, the lower and upper bounds of the IUP. The extension $\bar{\psi}^*(X_I)$ of this IUP is the set of UPs consistent with the bounds, i.e.,

$$\bar{\psi}^*(X_I) := \{ \psi : \Omega_{X_I} \rightarrow \mathbb{R} \mid \underline{\psi}(x_I) \leq \psi(x_I) \leq \bar{\psi}(x_I), \forall x_I \in \Omega_{X_I} \}. \quad (4)$$

The extension of a IUP $\bar{\psi}$ is non-empty if and only if $\underline{\psi}(x_I) \leq \bar{\psi}(x_I) \forall x_I \in \Omega_{X_I}$. We similarly define an *interval-valued probability potential* (IPP) over X_I given X_J as a pair of (not necessarily normalized) PPs over X_I given X_J . We denote such a IPP as $\bar{\phi}(X_I, X_J)$, where $\underline{\phi}(X_I, X_J)$ and $\bar{\phi}(X_I, X_J)$ are the two (unnormalized) bounds. The extension is also defined in terms of consistency.

$$\bar{\phi}^*(X) := \left\{ \phi : \Omega_{X_I} \times \Omega_{X_J} \rightarrow \mathbb{R}_0^+ \mid \begin{array}{l} \sum_{x_I} \phi(x_I, x_J) = 1, \\ \phi(x_I, x_J) \leq \bar{\phi}(x_I, x_J), \\ \forall (x_I, x_J) \in \Omega_{X_I} \times \Omega_{X_J} \end{array} \right\}. \quad (5)$$

Condition $\phi(x_I, x_J) \leq \bar{\phi}(x_I, x_J)$ for each x_I, x_J , together with $\sum_{x_I} \phi(x_I, x_J) \leq 1 \leq \sum_{x_I} \bar{\phi}(x_I, x_J)$, for each $x_J \in \Omega_{X_J}$ is necessary and sufficient for the extension of the IPP to be non-empty. The additional condition $\bar{\phi}(x'_i) + \sum_{x_i \neq x'_i} \phi(x_i) \leq 1$ and the analogous expression for the lower instead of the upper bounds is called *reachability* [5]. The meaning is that for each $p \in [\underline{\phi}(x_I, x_J), \bar{\phi}(x_I, x_J)]$, there is at least a PP $\phi \in \bar{\phi}^*$ s.t. $\phi(x_I, x_J) = p$. Note also that an IPP with non-empty extension can be always reduced to a reachable one by shrinking its bounds. Both the extensions of a IUP and a IPP are therefore convex sets of, respectively, UPs and PPs. Vice versa, while any convex set of UPs can be regarded as the extension of a IUP, the same does not hold for IPPs.

Example 3. Fig. 2 reports an interval-valued specification of the five potential associated to the ID in Fig. 1. It is a trivial exercise to check that: the IPPs have non-empty extensions and are reachable, and the UPs and PPs of the original model are included in the extensions of their interval-valued counterparts.

$\phi(O)$	$\psi(O, D)$	$\psi(T)$
$\begin{array}{c ccc} & e & w & s \\ \hline & [4875, 5125] & [2925, 3175] & [195, 22] \end{array}$	$\begin{array}{c ccc} & e & w & s \\ \hline d & [-75, -65] & [45, 55] & [195, 205] \\ T & [-5, 5] & [-5, 5] & [-5, 5] \\ \hline nd & & & \end{array}$	$\begin{array}{c cc} & t & nt \\ \hline & [-15, -5] & [-5, 5] \end{array}$
$\phi(S O, T=t)$	$\phi(S O, T=nt)$	
$\begin{array}{c ccc} & e & w & s \\ \hline c & [.0975, .1225] & [.2925, .3175] & [.4875, .5125] \\ S & [.2925, .3175] & [.39, .415] & [.39, .415] \\ d & [.585, .61] & [.2925, .3175] & [.0975, .1225] \end{array}$	$\begin{array}{c ccc} & e & w & s \\ \hline c & [.325, .35] & [.325, .35] & [.325, .35] \\ S & [.325, .35] & [.325, .35] & [.325, .35] \\ d & [.325, .35] & [.325, .35] & [.325, .35] \end{array}$	

Fig. 2. A set of IUPs and IPPs for the oil wildcatter's decision problem

Combining Interval-Valued Potentials. The combination operation over potentials in Alg. 1 can be extended to interval-valued potentials as follows:

- i) given two IUPs, say $\bar{\psi}(X_I)$ and $\bar{\psi}'(X_J)$, their combination $\bar{\psi} \otimes \bar{\psi}'$ is a IUP over $X_{I \cup J}$ s.t. $(\bar{\psi} \otimes \bar{\psi}')(x_{I \cup J}) := \bar{\psi}(x_I) + \bar{\psi}'(x_J)$ for each $x_{I \cup J} \in \Omega_{X_{I \cup J}}$ with $x_I, x_J \sim x_{I \cup J}$; and similarly for the upper bounds;
- ii) given a IPP $\bar{\phi}(X_I, X_J)$ and a IUP $\bar{\psi}(X_K)$, their combination $\bar{\phi} \otimes \bar{\psi}$ is a IUP over $X_{I \cup J \cup K}$ s.t. $(\bar{\phi} \otimes \bar{\psi})(x_{I \cup J \cup K}) := \bar{\phi}(x_I, x_J) \cdot \bar{\psi}(x_K)$ for each $x_{I \cup J \cup K} \in \Omega_{X_{I \cup J \cup K}}$, with $x_I, x_J, x_K \sim x_{I \cup J \cup K}$; if $\bar{\psi}(x_K) < 0$ the lower bound of the combination is obtained by multiplying the lower bound of the IUP for the upper bound of the IPP (and vice versa for the upper bound);
- iii) given two IPPs, say $\bar{\psi}(X_I, X_J)$ and $\bar{\psi}'(X_K, X_L)$, their combination $\bar{\psi} \otimes \bar{\psi}'$ is a IPP over $X_{I \cup K}$ given $X_{(J \cup L) \setminus (I \cup K)}$ s.t. $(\bar{\psi} \otimes \bar{\psi}')(x_{I \cup K}, x_{(J \cup L) \setminus (I \cup K)}) := \bar{\psi}(x_I, x_J) \cdot \bar{\psi}'(x_K, x_L)$, for each $x_{I \cup K} \in \Omega_{X_{I \cup K}}$ and $x_{(J \cup L) \setminus (I \cup K)} \in \Omega_{X_{(J \cup L) \setminus (I \cup K)}}$, with $x_I, x_J, x_K, x_L \sim x_{I \cup K}, x_{(J \cup L) \setminus (I \cup K)}$. If $\bar{\psi} \otimes \bar{\psi}'$ is not reachable, the transformation to make it reachable is performed.

The following result, whose proof is left to the reader provides a sensitivity-analysis justification for the proposed generalization of the combination operator.

Proposition 1. *Given potentials (no matter whether IUPs or IPPs) $\underline{\psi}$ and $\underline{\phi}$ the extension of their combination is s.t. $(\underline{\psi} \otimes \underline{\phi})^* = \left\{ \psi \otimes \phi \mid \psi \in \underline{\psi}^*, \phi \in \underline{\phi}^* \right\}$.*

4 Interval Influence Diagrams

IDs can be extended to the interval framework by simply replacing the PPs and UPs in Def. 1 with an equal number of IPPs and IUPs defined on the same domains. We call a model of this kind an *interval-valued influence diagram* (IID). As an example, the ID in Ex. 1 with the interval-valued quantification in Fig. 2 is an IID over the graph in Fig. 1. Before applying to IIDs the VE scheme in Alg. 1, the different operations over the potentials involved in the algorithm should be extended to intervals. In the previous section we discussed how to do that with the combination. Here we generalize the operations in lines 4 and 6.

Eliminating Chance Variables. We start from line 4. To sum out a variable from a IPP, we sum out the variable from the two bounds and, if needed, we make the result reachable. Concerning the sum and division in the second term, given an IPP $\underline{\phi}(X_I \cup Y, X_J)$ and a IUP $\underline{\psi}(X_K \cup Y)$, we set the result as a IUP $\underline{\hat{\psi}}$ over $X_{I \cup J \cup K}$ s.t.

$$\underline{\hat{\psi}}^*(X_{I \cup J \cup K}) := \left\{ \hat{\psi}(X_{I \cup J \cup K}) \left| \begin{array}{l} \hat{\psi}(x_{I \cup J \cup K}) = \frac{\sum_y \phi(x_I, y, x_J) \cdot \psi(y, x_K)}{\sum_y \phi(x_I, y, x_J)} \\ \forall x_{I \cup J \cup K}, \forall \phi \in \underline{\phi}^*, \forall \psi \in \underline{\psi}^* \end{array} \right. \right\}. \quad (6)$$

This allows to obtain a result as in Prop. 1. Note that the argument of the sum in the numerator can be regarded as an element of $(\underline{\phi} \otimes \underline{\psi})^*$. By computing the bound of the extension in Eq. (6), we eventually obtain the required IUP. Because of Eqs. (4) and (5), the extensions $\underline{\psi}^*$ and $\underline{\phi}^*$ are convex regions defined by linear constraints. Furthermore, the objective function to optimize is linear-fractional. So the task reduces to a linear program by the Charnes-Cooper transformation. Equivalently, the task can be regarded as a combinatorial optimization by considering only the extreme points of the feasible region.⁴ Faster, but approximate, approaches can be also considered. To obtain the lower bound, we can take the lower bound of the numerator and the upper bound of the denominator in Eq. (6) (and vice versa for the upper bound). This induces an outer approximation. A heuristic alternative consists in consider as PP specifications a lower bound for a value of Y and the upper bounds for the other values (or vice versa).

Eliminating Decision Variables. Here we discuss the operations in lines 6 and 7. The arg max operation is intrinsically related to the fact that a UP has sharp values. The problem of deciding the “maximal” options in the interval case might be arguable. The most conservative approach is the following.

Definition 3. *Let $\underline{\psi}$ be a IUP over $Y \cup \mathbf{X}_I$. An element $y \in \Omega_Y$ is interval-maximal given $x_I \in \Omega_{\mathbf{X}_I}$ if there is no $y' \in \Omega_Y$ s.t. $\underline{\psi}(y', x_I) > \underline{\psi}(y, x_I)$.*

⁴ The solution of a linear program is an extreme point of the feasible region.

Let D be a decision to be eliminated from $\overline{\psi}(D, \mathbf{X}_I)$. To detect the optimal policy $\delta_D^*(\mathbf{X}_I)$ we compute the interval-maximal states of D given each $x_I \in \Omega_{\mathbf{X}_I}$. This corresponds to a so-called *credal* policy allowing for indecision between two or more possible options. Finally the maximization of the IUP is done as usual by acting separately on the two bounds.

Example 4. In the oil wildcutter’s problem (Ex. 1), if the combinatorial approach is used, $\delta_D^*(S = c, T = t) = \{d\}$ as $EU_D(S = c, T = t) \in [46.24, 115.78]$.

Complexity Analysis. VE in IDs takes time exponential in the maximum clique size (i.e., the arity of the biggest potential generated during the evaluation). The same holds for IIDs, apart from a possible bottleneck during the elimination of the chance variables as in Eq. (6). This is the case when the combinatorial optimization is adopted: the number of extreme points to be evaluated is exponential in the number of states. This is not the case with the outer approximation and the heuristic, as well as the linear programming which is polynomial in the number of constraints, and hence in the number of states.

5 Empirical Validation

For a preliminary validation of the VE algorithm we consider six IDs [1, 3, 13, 15]. Tab. 1 details the number of nodes of each type for these models. These IDs are transformed in IIDs by a perturbation of the original parameters. The approaches in Sect. 4) are compared. Fig. 3 shows the computation times. As expected, the outer approximation and the heuristic method roughly take the double of the time required by the precise evaluation. The exact approach with the extreme points might be very slow if there are chance variables with many states, this being the case of NHL, which has a chance variable with 12 states. The exact approach based on linear programming is slightly slower than the heuristic and the outer approximation.

	NHL	Jaundice	Appendicitis	Comp. Assym	Oil	Thinkbox
$ \mathbf{X} $	17	21	4	3	2	5
$ \mathbf{D} $	3	2	1	5	2	2
$ \mathbf{U} $	1	1	1	2	1	4

Table 1. Number of chance, decision and utility nodes for the benchmark IIDs

A sensitivity analysis has also been done to evaluate the effect of the size of the intervals of the initial potentials affects the informativeness of the solutions. Fig. 4 shows the size of the (interval-valued) expected utility of the optimal policy $EU(\Delta^*)$ as a function of the sizes of the intervals of the potentials. The results based on the linear programming and on the enumeration of the extreme points are exact and therefore coincide. As expected the heuristic is more precise than the outer approximation, which in some cases returns infinite values because of a division by zero. If sharp values are assumed for the IPPs and the intervals are only in the IUPs, the results are those in Fig. 5. The four methods, which differs only in the treatment of the IPPs, produce the same results. Comparing the scales of Figs. 4 and 5, it can be observed that the imprecision in the IPPs has a stronger effect than in the IUPs.

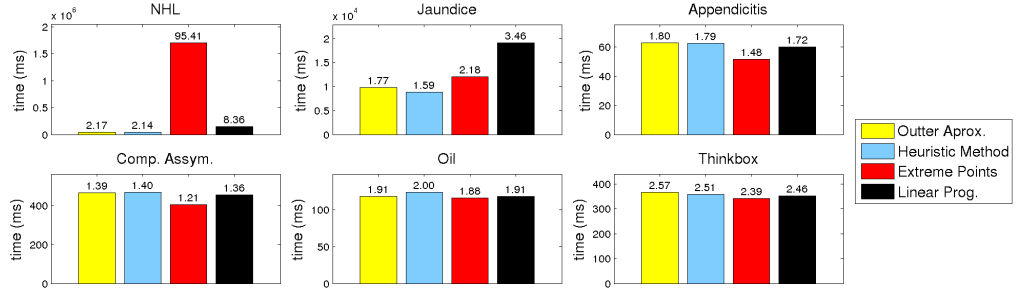


Fig. 3. Evaluation times for the IIDs in Tab. 1 and relative duration w.r.t. the evaluation of the corresponding IDs

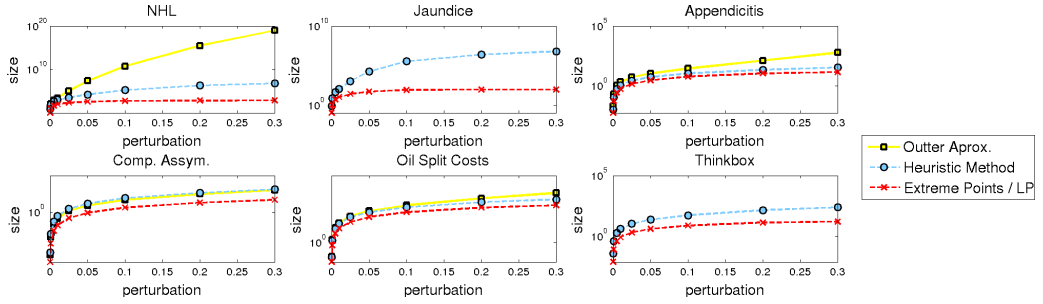


Fig. 4. Size of $EU(\Delta^*)$ for different sizes of the IPPs

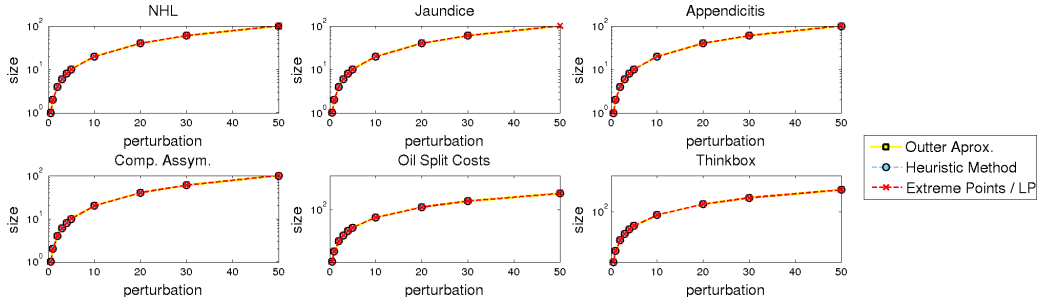


Fig. 5. Size of $EU(\Delta^*)$ for different sizes for the IUPs

6 Conclusions and Future Work

We have generalized the formalism of influence diagrams to the interval framework by allowing both probabilities and utilities to take interval values. A variable elimination algorithm has been also proposed and preliminary tested. In the experimental part, four different methods for eliminating chance variables have been compared, showing that the best results are obtained if the linear programming approach is considered. As a future work intend to develop approximate evaluation algorithms for these models.

Acknowledgments

This research was supported by the Spanish Ministry of Economy and Competitiveness under project TIN2013-46638-C3-2-P, the European Regional Development Fund (FEDER), the FPI scholarship program (BES-2011-050604) and the short stay in foreign institutions scholarship EEBB-I-14-08102. The authors have also been partially supported by “Junta de Andalucía” under projects TIC-06016 and P08-TIC-03717.

References

1. Hugin Expert network repository. <http://www.hugin.com/technology/samples>.
2. S. Benferhat and S. Smaoui. Hybrid possibilistic networks. *International Journal of Approximate Reasoning*, 44(3):224 – 243, 2007.
3. C. Bielza, M. Gómez, Ríos S. Insua, J. A. Fernández del Pozo, Barreno García P., S Caballero, and M. Sánchez Luna. IctNEO system for jaundice management. *Revista de la Real Academia de Ciencias Exactas, Físicas y Naturales*, 92(4):307–315, 1998.
4. F. G. Cozman. Credal networks. *Artificial Intelligence*, 120:199–233, 2000.
5. L. M. de Campos, J. F. Huete, and S. Moral. Probability intervals: a tool for uncertain reasoning. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 2(02):167–196, 1994.
6. E. Fagioli and M. Zaffalon. Decisions under uncertainty with credal influence diagrams. Technical Report 51-98, IDSIA, 1998. (unpublished).
7. K. W. Fertig and J. S. Breese. Probability intervals over influence diagrams. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(3):280–286, 1993.
8. R. A. Howard and J. E. Matheson. Influence diagram retrospective. *Decision Analysis*, 2(3):144–147, 2005.
9. N. Huntley and M. C. M. Troffaes. Normal form backward induction for decision trees with coherent lower previsions. *Annals of Operations Research*, 195(1):111–134, 2012.
10. F.V. Jensen and T.D. Nielsen. *Bayesian networks and decision graphs*. Springer Verlag, 2007.
11. D. Kikuti, F. G. Cozman, and C. P. de Campos. Partially ordered preferences in decision trees: computing strategies with imprecision in probabilities. In *IJCAI workshop on advances in preference handling*, pages 118–123, 2005.
12. U. Kjaerulff. Triangulation of graphs – algorithms giving small total state space. Research Report R-90-09, Department of Mathematics and Computer Science, Aalborg University, Denmark, 1990.
13. P.J.F. Lucas and B. Taal. Computer-based decision support in the management of primary gastric non-hodgkin lymphoma. *UU-CS*, 33, 1998.
14. T.D. Nielsen and F.V. Jensen. Sensitivity analysis in influence diagrams. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 33(2):223–234, 2003.
15. H. Raiffa. *Decision analysis: Introductory lectures on choices under uncertainty*. 1968.
16. P. P. Shenoy. Valuation-based systems for Bayesian decision analysis. *Operations research*, 40(3):463–484, 1992.
17. H. Xu and P. Smets. Reasoning in evidential networks with conditional belief functions. *International Journal of Approximate Reasoning*, 14(2-3):155 – 185, 1996.
18. N.L. Zhang and D. Poole. Exploiting causal independence in Bayesian network inference. *Journal of Artificial Intelligence Research*, 5:301–328, 1996.