Cost-effectiveness analysis with influence diagrams

M. Arias and F. J. Díez

UNED, Madrid, Spain

Abstract. Cost-effectiveness analysis (CEA) is used more and more frequently in medicine to determine whether the health benefits of an intervention outweighs its economic cost. In this paper we present two algorithms for performing CEA with influence diagrams; one of them is based on variable elimination and the other on arc reversal. Using the former, we have performed CEA on two influence diagrams whose equivalent decision trees contain thousands of leaves.

Keywords: Cost-effectiveness analysis, net monetary benefit, influence diagrams.

1 Introduction

In medicine, one of the methods of assessing whether the health benefits of an intervention outweighs its economic cost is cost-effectiveness analysis (CEA) [4, 5]. In this context, the *net monetary benefit* [16] of an intervention I_i is

$$NMB_{I_i}(\lambda) = \lambda \cdot e_i - c_i \ , \tag{1}$$

where e_i is its effectiveness and c_i its cost. The parameter λ is used to convert the effectiveness into a monetary scale. It takes values on the set of positive real numbers, i.e., in the interval $(0, +\infty)$. It is measured in effectiveness units divided by cost units; for example, in dollars per death avoided or euros per quality adjusted life-year (QALY [19]). It is sometimes called *willingness to pay*, *cost-effectiveness threshold* or *ceiling ratio*, because it indicates how much money a decision maker accepts to pay to obtain a certain "amount" of health benefit.

When the consequences of an intervention are not deterministic, it is necessary to apply a model that takes into account the probability of each outcome. The most usual tool for modeling decision problems with uncertainty are decision trees [13]. In a previous paper [1] we have presented a method of performing CEAs on decision trees with an arbitrary number of decision nodes.

The main drawbacks of decision trees is that the size of a tree grows exponentially with the number of variables, they cannot represent conditional independencies, and they require a preprocessing of the probabilities [2, 6]; for instance, medical diagnosis problems are usually stated in terms of direct probabilities, namely the prevalence of the diseases and the sensitivity and specificity of the tests, while the tree is built with the inverse probabilities, i.e., the positive and negative predictive values of the tests. Even in cases with only a few chance variables, this preprocessing of probabilities is difficult, if not impossible.

An alternative representation language for decision making are influence diagrams (IDs) [6]. They have the advantages of being very compact, representing conditional independencies, and using direct probabilities. However, the only algorithm that can perform CEA with IDs is that proposed by Nielsen et al. [11], which is very difficult to apply in practice for the reasons discussed below. In this paper, we present two efficient algorithm for CEA with IDs that have allowed us to solve medical problems that were impossible to address with the techniques available so far. One of them is based on the variable elimination algorithm [7]; the other, on the arc reversal algorithm [12, 14].

The rest of this paper is structured as follows: Section 2 reviews the basic concepts of CEA and IDs. The new algorithms are presented in Section 3. Section 4 shows an example, Section 5 discusses some related work, and Section 6 contains the conclusions.

2 Background

2.1 Cost-effectiveness analysis (CEA)

Deterministic CEA Cost-effectiveness analysis (CEA) consists of finding the intervention that maximizes the net benefit for each value of λ (cf. Eq. 1). When we have a set of interventions such that the cost and effectiveness of each one are known with certainty, we can perform a deterministic cost-effectiveness analysis, which returns the optimal intervention for each interval of the possible values of λ. The standard algorithm for this analysis consists of eliminating the interventions dominated by another intervention (simple dominance), then eliminating the interventions dominated by a pair of other interventions (extended dominance), and finally computing the incremental cost-effectiveness ratios—see [18] or any book on medical decision analysis. This algorithm and an alternative method for deterministic CEA can be found in [1].

CEA with decision trees Sometimes we do not know explicitly the cost and effectiveness of each intervention, but we do know that each one may lead to different outcomes with different probabilities, which may in turn cause other outcomes, each having a known cost and effectiveness. In this case, we can build a decision tree such that each node, instead of representing a single utility, represents the cost and effectiveness of the corresponding scenario.

If the only decision node of the tree is the root, the tree can be evaluated by a modified version of the roll-back algorithm that computes the cost and effectiveness of each node separately and then performs a deterministic CEA at the root node [4, 15]. If the tree contains embedded nodes, this method cannot be applied because the evaluation of a decision node does not return a cost-effectiveness pair. However such a tree can be evaluated with the algorithm proposed in [1].

2.2 Influence diagrams

Basic properties of IDs An ID is a probabilistic graphical model that consists of three disjoint sets of nodes: decision nodes $\mathbf{V}_D$, chance nodes $\mathbf{V}_C$, and utility nodes $\mathbf{V}_U$. Chance nodes represent events that are not under the direct control of the decision maker. Decision nodes correspond to actions under the direct control of the decision maker. Given that each chance or decision node represents a variable, we will use indifferently the terms variable and node. Standard IDs require that there is a total ordering of the decisions, which indicates the order in which the decisions are made.

In this section, we assume that the ID contains only one utility node. However, the algorithms presented here can be easily extended to IDs in which the global utility is the sum of the values represented by the utility nodes [7] and to IDs containing *super-value* nodes [10, 17]. (A utility node is said to be *super-value* if its parents are other utility nodes.)

The meaning of an arc in an ID depends on the type of nodes that it links. An arc $X \to C$ where C is a chance node denotes a probabilistic dependence of C on X; in practice, it usually means that X is a cause of C. An arc from a decision D_i to a decision D_j means that D_i is made before D_j. An arc from a chance node C to a decision node D_j means that the value of variable C is known when making decision D_j. Standard IDs assume the *non-forgetting hypothesis*, which means that a variable C known for a decision D_j is also known for any posterior decision D_k, even if there is not an explicit link $C \to D_k$ in the graph. An arc from a variable X to the utility node means that the utility depends on X.

A *potential* is a real-valued function over a domain of finite variables. The quantitative information that defines an ID is given by assigning to each chance node C a conditional probability potential $P(c|pa(C))$ for each configuration of its parents, $pa(C)$ and assigning to the utility node a potential $U(pa(U))$ that maps each configuration of its parents onto a real number.

The total ordering of the decisions $\{D_1, \ldots, D_n\}$ induces a partition of the chance variables $\{\mathbf{C}_0, \mathbf{C}_1, \ldots, \mathbf{C}_n\}$, where $\mathbf{C}_i$ is the set of variables unknown for D_i and known for D_{i+1}. The set of variables known to the decision maker when deciding on D_i is called the *informational predecessors* of D_i and denoted by $\mathit{IPred}(D_i)$. Consequently,

$$\begin{aligned} \mathit{IPred}(D_i) &= \mathbf{C}_0 \cup \{D_1\} \cup \mathbf{C}_1 \cup \ldots \cup \{D_{i-1}\} \cup \mathbf{C}_{i-1} && (2) \\ &= \mathit{IPred}(D_{i-1}) \cup \{D_{i-1}\} \cup \mathbf{C}_{i-1} \; . && (3) \end{aligned}$$

The *maximum expected utility* (*MEU*) of an ID whose chance and decision variables are all discrete is

$$\mathit{MEU} = \sum_{\mathbf{c}_0} \max_{d_1} \sum_{\mathbf{c}_1} \ldots \sum_{\mathbf{c}_{n-1}} \max_{d_n} \sum_{\mathbf{c}_n} \prod_{C \in \mathbf{V}_C} P(c|pa(C)) \cdot U(pa(U)) \; . \quad (4)$$

A *policy* δ_{D_i} is a function that maps each configuration of informational predecessors of D_i onto a value d_i of D_i. The *optimal policy* δ_{D_i} for decision D_i

is given by the following equation (in the case of a tie, any of the values of D_i that maximize that expression can be chosen arbitrarily):

$$\delta_{D_i}(IPred(D_i)) = \arg\max_{d_i} \sum_{\mathbf{c}_i} \max_{d_{i+1}} \ldots \sum_{\mathbf{c}_{n-1}} \max_{d_n} \sum_{\mathbf{c}_n} \prod_{C \in \mathbf{V}_C} P(c|pa(C)) \cdot U(pa(U)) \ . \quad (5)$$

Algorithm 1: Variable elimination with division of potentials

Input: An influence diagram.
Result: The expected utility and the optimal policy for each decision.
`// Initialize the list of probability potentials`
1 *list* $\leftarrow \{P(c|pa(C)) \mid C \in \mathbf{V}_C\}$;
2 **for** $i \leftarrow n$ **to** 0 **do**
    `// eliminate the variables in` $\mathbf{C}_i$
3   **foreach** *variable* $C \in \mathbf{C}_i$ **do**
4   take out from *list* all the potentials that depend on C;
5 $\psi \leftarrow$ product of those potentials;
6 $\psi_{\text{ind}} \leftarrow \sum_c \psi$;
7 $\psi_{\text{cond}} \leftarrow \psi / \psi_{\text{ind}}$;
8   add ψ_{ind} to the *list*;
9 $U \leftarrow \sum_c \psi_{\text{cond}} \cdot U$;
10   **if** $i > 0$ **then**
        `// eliminate the decision` D_i
11 $U \leftarrow \max_{d_i} U$;
12 $\delta_{D_i}(IPred(D_i)) = \arg\max_{d_i \in D_i} U$;
13   **foreach** *potential* $P(c|pa(C)) \in$ *list* **do**
14   **if** *this potential depends on* D_i **then**
15   project this potential onto the remaining variables
16   in the *list*;
17   replace $P(c|pa(C))$ with its projection;
18 return U

Variable elimination The direct application of the above expression leads to a computational cost that grows exponentially with the number of variables in the ID. A more efficient approach consists of eliminating the variables one by one, in an order compatible with the above equations, i.e., eliminating first the variables that appear on right-most operators of summation and maximization in Equation 4, as indicated by the Algorithm 1. The details and the justification of this algorithm can be found in [7, 10].

Please note that the potential ψ, which is the product of all the potentials that depend on C (line 5), is factored into two potentials: ψ_{ind}, where "ind" stands for "independent of C", and ψ_{cond}, where c stands for "conditional probability" because this potential represents the conditional probability of C given the variables that have not been eliminated yet: $\psi_{\text{cond}} = P(c|\mathbf{v}_C)$. The meaning of this probability can be better understood if we consider that for each influence diagram there is an equivalent symmetric decision tree. The variables eliminated before C when evaluating the ID are placed on the right side of C in the decision tree, and those eliminated after C are placed on the left side. Each configuration $\mathbf{v}_C$ represents a path from the root node to a node C in the tree, and the probability $P(c|\mathbf{v}_C)$ computed in the evaluation of the ID is the probability of the branches outgoing from that node in the tree.

Then the algorithm multiples the potentials $\psi_{\text{cond}} = P(c|\mathbf{v}_C)$ and U, and sums out the variable C (line 9). This is equivalent to evaluating all the C nodes in the tree by computing the average of the utilities of their branches, using $P(c|\mathbf{v}_C)$ as the weights. The elimination of variable D_i by maximizing U (line 11) is equivalent to evaluating all the D_i nodes in the tree.

When all the chance and decision variables have been eliminated, the potential U contains only one real number, that is the expected utility of the ID (line 18).

Arc reversal An alternative method for evaluating IDs is the arc reversal algorithm proposed by Olmsted [12] (see also [14]). This algorithm consists of four basic operations:

1. **Barren node removal.** A node is said to be *barren* if it has no children. Barren nodes can be removed from the ID without performing any additional operation.
2. **Chance node removal.** A chance node Y whose only child is the utility node U can be removed from the ID by drawing links from each of its parents to U; if $\mathbf{X}$ is the set of parents of Y, $P_{\text{old}}(y|\mathbf{x})$ is the conditional probability of Y, $\mathbf{Z}$ is the set of parents of U, and $U_{\text{old}}(\mathbf{z})$ is the utility potential before eliminating Y, the new utility potential after eliminating Y is
$$U_{\text{new}}(\mathbf{v}) = \sum_y P_{\text{old}}(y|\mathbf{x}) \cdot U_{\text{old}}(\mathbf{z}) \ , \tag{6}$$
where $\mathbf{V} = \mathbf{X} \cup (\mathbf{Z} \backslash \{Y\})$.
3. **Decision node removal.** A decision node D whose only child is the utility node U can be removed from the ID by drawing links from each of its parents to U; if $\mathbf{X}$ is the set of parents of D, $\mathbf{Z}$ is the set of parents of U, and $U_{\text{old}}(\mathbf{z})$ is the utility potential before eliminating D, the new utility potential after eliminating D is
$$U_{\text{new}}(\mathbf{v}) = \max_d U_{\text{old}}(\mathbf{z}) \ , \tag{7}$$
where $\mathbf{V} = \mathbf{X} \cup (\mathbf{Z} \backslash \{D\})$.
4. **Arc reversal.** A link $X \rightarrow Y$ in the ID such that there is no other directed from X to Y in the graph can be reversed, i.e., replaced by $Y \rightarrow X$, by performing the following additional operations: (1) all the parents of X become

parents of Y, and vice versa, (2) the new probability of Y is

$$P_{\text{new}}(y|\mathbf{a},\mathbf{b},\mathbf{c}) = \sum_{x} P_{\text{old}}(x|\mathbf{a},\mathbf{b}) \cdot P_{\text{old}}(y|\mathbf{b},\mathbf{c}) \ , \tag{8}$$

and the new probability of X is

$$P_{\text{new}}(x|y,\mathbf{a},\mathbf{b},\mathbf{c}) = \sum_{y} \frac{P_{\text{old}}(x|\mathbf{a},\mathbf{b}) \cdot P_{\text{old}}(y|\mathbf{b},\mathbf{c})}{P_{\text{new}}(y|\mathbf{a},\mathbf{b},\mathbf{c})} \ , \tag{9}$$

where $\mathbf{A} = Pa(X) \backslash Pa(Y)$, $\mathbf{B} = Pa(X) \cap Pa(Y)$, and $\mathbf{C} = Pa(Y) \backslash \{Pa(X) \cup X\}$.

Each of these operations transforms an ID into an equivalent ID having the same optimal policies and the same expected utility (except for the decisions removed, obviously). The Theorem 4 in [14] states that if the ID does not contain any node that can be removed directly, there exists a sequence of arc reversals leading to an equivalent ID in which at least one node can be removed. The elimination of a decision node D gives the optimal policy for D. The utility potential remaining after eliminating all the chance and decision nodes contains a single real number (a scalar), that is the expected utility of the ID. The algorithm always terminates because the number of nodes in the ID is finite.

3 Cost-effectiveness analysis with IDs

3.1 Cost-effectiveness partitions

As mentioned above, cost-effectiveness analysis consists of finding an intervention that maximizes the net benefit for each value of λ; in practice, it consists of finding the intervals for which an intervention is more beneficial than the others. We formalize this idea by introducing the concept of *cost-effectiveness partition*, CEP.

Definition 1. *A* cost-effectiveness partition *(CEP) of n intervals is a tuple $Q = (\Theta_Q, C_Q, E_Q, I_Q)$, where:*

- $\Theta_Q = \{\theta_1, \ldots, \theta_{n-1}\}$ is a set of $n-1$ values (thresholds), such that $0 < \theta_1 < \ldots < \theta_{n-1}$,
- $C_Q = \{c_0, \ldots, c_{n-1}\}$ is a set of n values (costs),
- $E_Q = \{e_0, \ldots, e_{n-1}\}$ is a set of n effectiveness values, and
- $I_Q = \{I_0, \ldots, I_{n-1}\}$ is a set of n interventions.

For the sake of simplifying the exposition, we define $\theta_0 = 0$ and $\theta_n = +\infty$ for every CEP.

Alternatively, a CEP can be denoted by a set of n 4-tuples of the form (interval, cost, effectiveness, intervention),

$$\begin{aligned} Q = \{ & ((0, \theta_1), c_0, e_0, I_0), \\ & ((\theta_1, \theta_2), c_1, e_1, I_1), \\ & \ldots, \\ & ((\theta_{n-1}, +\infty), c_{n-1}, e_{n-1}, I_{n-1})\} , \end{aligned}$$

which means that when λ is in the interval (θ_i, θ_{i+1}) the most beneficial intervention is I_i, which has a cost c_i and an effectiveness c_i. When $\lambda = \theta_{i+1}$, there is a tie between I_i and I_{i+1}.

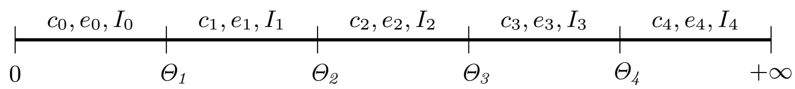

Fig. 1. Cost-effectiveness partition (CEP).

The functions $cost_Q(\lambda)$, $eff_Q(\lambda)$, $NMB_Q(\lambda)$, and $interv_Q(\lambda)$ return the cost, the effectiveness, the NMB, and the optimal intervention for λ according to the CEP Q; see [1] for a mathematical definition of these functions.

Combination of cost-effectiveness partitions In this subsection we generalize the average and maximization operations from single utilities to CEPs.

Definition 2 (Weighted average). *Given a set of m CEPs $\{Q_1, \ldots, Q_m\}$, a chance variable X whose domain is $\{x_1, \ldots, x_m\}$, and a probability distribution for X, $P(x_j)$, we say that a CEP Q is a* weighted average *of the CEPs if*

$$\forall \lambda, \quad cost_Q(\lambda) = \sum_{j=1}^{m} P(x_j) \cdot cost_{Q_j}(\lambda) \tag{10}$$

and

$$\forall \lambda, \quad eff_Q(\lambda) = \sum_{j=1}^{m} P(x_j) \cdot eff_{Q_j}(\lambda) \ . \tag{11}$$

A straightforward consequence of this definition is that, because of Equation 1,

$$\forall \lambda, \quad NMB_Q(\lambda) = \sum_{j=1}^{m} P(x_j) \cdot NMB_{Q_j}(\lambda) \ . \tag{12}$$

These three equalities mean that for every value of λ, the cost, effectiveness, and NMB of the weighted average partition Q are the same as if we had performed a weighted average for the values of cost and effectiveness of the Q_j's. The partition Q can be efficiently computed by the Algorithm 2.

The intervention composed at the fifth line of the algorithm means: "if the chance variable X takes on the value x_j, then follow the policy indicated by the corresponding branch $(X = x_i)$ of the tree".

Algorithm 2: Weighted average of CEPs

Input: A set of m CEPs $\{Q_1, \ldots, Q_m\}$, with $Q_j = (\Theta_j, C_j, E_j, I_j)$,
a chance variable X whose domain is $\{x_1, \ldots, x_m\}$, and
a probability distribution for X, $P(x_j)$.

Result: A new CEP $Q = (\Theta, C, E, I)$.

1 $\Theta \leftarrow \bigcup_{j=1}^{m} \Theta_j$
2 $n \leftarrow \text{card}(\Theta)$
3 **for** $i \leftarrow 1$ **to** n **do**
4     $c_i \leftarrow \sum_{j=1}^{m} P(x_j) \cdot cost_{Q_j}(\theta_i)$
5     $e_i \leftarrow \sum_{j=1}^{m} P(x_j) \cdot e\!f\!f_{Q_j}(\theta_i)$
6     $I_i \leftarrow$ "If $X = x_1$, then $interv_{Q_1(\theta_i)}$; if $X = x_2$, then $interv_{Q_2(\theta_i)}$; etc."

Definition 3 (Optimal partition). *Given a set of m CEPs $\{Q_1, \ldots, Q_m\}$ and a decision D whose domain is $\{d_1, \ldots, d_m\}$, a CEP Q is* optimal *if*

$$\forall \lambda,\ \exists j,\ \ NMB_{interv_{Q_j}(\lambda)}(\lambda) = max_{j'} NMB_{interv_{Q_{j'}}(\lambda)}(\lambda)\ , \tag{13}$$

$$interv_Q(\lambda) = \text{"choose option } d_j\text{; then apply } interv_{Q_j}(\lambda)\text{"}\ , \tag{14}$$

$$cost_Q(\lambda) = cost_{Q_j}(\lambda)\ , \tag{15}$$

$$e\!f\!f_Q(\lambda) = e\!f\!f_{Q_j}(\lambda)\ . \tag{16}$$

The interpretation of this definition is as follows: for each value d_j (a possible choice) of the decision D there is CEP Q_j and for each value of λ there is an intervention $interv_{Q_j}(\lambda)$ in Q_j which is optimal for d_i. Equation 13 means that when making decision D we select j such that $interv_{Q_j}(\lambda)$ is the intervention that attains the highest NMB for that particular value of λ. Equation 14 means that the optimal intervention for decision D is to choose first the option d_j and then apply the intervention $interv_{Q_j}(\lambda)$. The cost and effectiveness associated with intervention $interv_Q(\lambda)$—given by the optimal CEP, Q—are the same as in Q_j.

The key property of this definition is Equation 13, which states that for every λ the NMB of the optimal partition, Q, is the same as if we had performed a (unicriterion) maximization of the NMB for each single value of λ.

The optimal CEP can be obtained by applying Algorithm 3, which collects all the thresholds of the Q_j's and performs a deterministic CEA (cf. Sec. 2.1) on each interval. Finally, it fuses some intervals by eliminating the unnecessary thresholds. In [1] we show with an example how this algorithm operates and why it is sometimes necessary to fuse intervals.

Algorithm 3: Optimal CEP.

Input: A set of m CEPs $\{Q_1, \ldots, Q_m\}$, with $Q_j = (\Theta_j, C_j, E_j, I_j)$ and a decision node
Result: A new CEP $Q = (\Theta, C, E, I)$.

$\Theta \leftarrow \bigcup_{j=1}^{m} \theta_j$
$n \leftarrow$ card(Θ)
for $i \leftarrow 1$ **to** n **do**
  perform a deterministic CEA analysis of interval i
fuse contiguous intervals having the same intervention, the same cost, and the same effectiveness

3.2 Construction of the ID

The construction of an ID for CEA is almost identical to the traditional case, in which the ID contains only one utility node because the decision is based on a single criterion. The difference is that in CEA we have two criteria, and consequently we put two utility nodes in the ID: U_c for the cost and U_e for the effectiveness. We will see an example below.

3.3 Evaluation of the influence diagram

Evaluation of the ID with the variable elimination algorithm Performing CEA with an ID is very similar to the evaluation of a (unicriterion) ID having only one decision node. The main differences are:

- In the unicriterion case, each potential assigns a real number to each configuration of its variables, while now we have a CEP-potential that assigns a CEP to each configuration.
- Initially, U is a CEP-potential that depends on all the variables that are parents of U_c or U_e. The CEP assigned to each configuration contains only one interval (no thresholds have been generated yet); the cost and the effectiveness are those obtained from the functions associated to U_c and U_e, which are kept separately in the CEP-potential.
- The weighted average and the maximization of potentials (lines 9 and 11 of the Algorithm 1) must be replaced by the weighted average of CEPs (Algorithm 2) and the computation of the optimal CEP (Algorithm 3).
- The potential returned by the algorithm is not a single real number, but a CEP.

The mathematical ground for this algorithm is as follows. If we know the value of λ, we can transform the cost-effectiveness problem into a one criterion problem by computing the NMB of each scenario using Equation 1. This way, the ID would have only one utility node instead of two and we might evaluate it with the Algorithm 1. As mentioned in the introduction, CEA is performed when we do not know the value of λ. The algorithm presented in this subsection is equivalent to applying the Algorithm 1 for every single value of λ, but instead of doing an independent evaluation for each value of λ, which is clearly impossible, we group the λ's into intervals having the same cost, the same effectiveness, and the same optimal intervention, and we evaluate all the values of λ in parallel. The CEP returned by the modified algorithm indicates, for each interval, the cost, the effectiveness, and the optimal intervention.

Evaluation of the ID with the arc reversal algorithm The basic idea of our method can be applied to the arc-reversal algorithm described in Section 2.2. The first step is to fuse the two utility nodes, U_c and U_e, into a single node U having an associated CEP-potential, the same as in the variable elimination algorithm for CEA. The operations of barren node removal and arc reversal are exactly the same as in the evaluation of unicriterion IDs.

The removal of a chance node is analogous to the case of unicriterion IDs, but now we have to perform a weighted average of CEPs (Algorithm 2) for each configuration of $\mathbf{V}$ (cf. Eq. 6), because now $U_{\text{old}}(\mathbf{z})$ is a CEP-potential. Similarly, when removing a decision node we have to find the optimal CEP (Algorithm 3) for each configuration of $\mathbf{V}$ (cf. Eq. 7).

4 Example: CEA of a test

Example 1. For a disease whose prevalence is 0.14, there are two possible therapies. The effectiveness of each therapy depends on whether the disease is present or not, as shown in Table 1. There is a test with a sensitivity of 90% and a specificity of 93%, and a cost of 150 €. Is the test cost-effective?

Therapy	**Cost**	**Effectiveness** +disease	**Effectiveness** ¬disease
No therapy	0 €	1.2	10,0
Therapy 1	20,000 €	4.0	9.9
Therapy 2	70,000 €	6.5	9.3

Table 1. Cost and effectiveness of each intervention for the Example 1.

This problem can be analyzed with the ID in Figure 2. The decision node *Dec:Test* represents the decision of performing the test or not, and *Therapy* represents the choice of therapy. The numerical information of the ID consists of four ordinary potentials (by "ordinary" we mean that they assign a real number to each configuration of their variables): $P(disease)$, $P(test|disease,dec{:}test)$, $U_e(disease,therapy)$, and $U_c(dec{:}test,therapy)$.

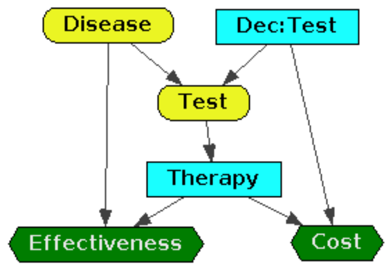

Fig. 2. Influence diagram for the Example 1.

Evaluation of this example with variable elimination The algorithm is initialized by building a CEP-potential U that depends on *Disease*, *Dec:Test*, and *Therapy*. For each of the $2 \times 2 \times 3 = 12$ configurations of these variables the CEP-potential U contains a CEP of only one interval, $(0, +\infty)$.

The first variable that the algorithm eliminates is *Disease*. The potential ψ is computed as the product of $P(disease)$ and $P(test|disease,dec\!:\!test)$. The potentials $\psi_{\text{ind}}(dec\!:\!test,test)$ and $\psi_{\text{cond}}(disease,dec\!:\!test,test) = P(disease|dec\!:\!test,test)$ are computed as indicated by the lines 6 and 7 in the Algorithm 1. The new CEP-potential U is

$$U(dec\!:\!test,test,therapy) = \sum_{disease} P(disease|dec\!:\!test,test) \cdot U(disease,dec\!:\!test,therapy) \ .$$

The elimination of the decision node *Therapy* is performed by applying a deterministic CEA (see Sec. 2.1) for each CEP in the new CEP-potential $U(dec\!:\!test,test,therapy)$. This returns a new CEP-potential, $U(dec\!:\!test,test)$.

The elimination of the chance node *Test* gives the following CEP-potential:

$$U(dec\!:\!test) = \sum_{test} \psi(dec\!:\!test,test) \cdot U(dec\!:\!test,test) \ .$$

Finally, the elimination of *Dec:Test* performs a deterministic CEA for each CEP in this CEP-potential, $U(dec\!:\!test)$, which returns the CEP shown in Table 2.

Evaluation of this example with arc reversal The first step consists of fusing the two utility nodes, *Effectiveness* and *Cost*, into a single node U with the same associated CEP-potential as in the case of variable elimination: $U(disease,dec\!:\!test,therapy)$.

The inversion of the arc *Disease*→*Test* leads to computing the potentials $P_{\text{new}}(disease|test,dec\!:\!test)$ and $P_{\text{new}}(test|dec\!:\!test)$. The first one is lost when eliminating the node *Disease*. The second is exactly the same potential $\psi(dec\!:\!test,test)$ computed by the variable elimination algorithm. The consecutive elimination of

Interval	Cost	Effectiveness	Dec:Test	Therapy
(0, 11,171)	0	8.77	Do not test	No therapy
(11,171, 33,384)	3,874	9.11	Do test	{ test:positive→Therapy 1 test:negative→No therapy
(33,384, $+\infty$)	13,184	9.39	Do test	{ test:positive→Therapy 2 test:negative→No therapy

Table 2. Final CEP obtained by evaluating the influence diagram in Figure 2. It gives the cost, the effectiveness, and the optimal intervention for each value of λ.

Therapy, *Test*, and *Dec:Test*, which do not require any arc reversal, lead to the same utility potentials as in the case of variable elimination. In particular, the final CEP is the one shown in Table 2.

This example shows that even though the two algorithms look different, in fact they are performing essentially the same operations.

5 Related work

Nielsen et al. [11] have studied multi-attribute IDs, in which the global utility is given by

$$u = \alpha_1 \cdot u_1 + \ldots + \alpha_n \cdot u_n \tag{17}$$

where each u_i is an attribute and the α_i's represent the decision maker's preferences. The IDs presented in this paper are a particular case of the former, in which $n = 2$. However, the evaluation methods are completely different. The work by Nielsen et al. focuses on a particular configuration α, assumes that $\mathbf{\Delta}^*_\alpha$, the optimal strategy for this α, is known (it is easy to evaluate a unicriterion ID), and tries to determine the support for this strategy, i.e., the region of $\mathbb{R}^n$ for which the optimal strategy is $\mathbf{\Delta}^*_\alpha$. The result of this analysis is a set of inequalities, which can be interpreted as the hyperplanes that delimit such region. A more intuitive way of summarizing the results of their analysis is to compute the radius of the biggest ball that can be contained in that region. In our opinion, this kind of analysis is far from the CEA that health decision makers demand.

In contrast, our study is limited to the case $n = 2$ and $\alpha_2 = -1$ (see Eq. 1), which allows us to find the optimal strategy as a function of α_1, i.e., λ. We perform the kind of CEA that is usual in medicine, and the output of our algorithm can be summarized in the form of a table (see Table 2) whose interpretation is immediate: it shows the cost, the effectiveness, and the optimal intervention for each value of λ.

6 Conclusion and future work

In this paper we have presented two algorithms for performing CEA with influence diagrams (IDs); one of them is based on the variables elimination algorithm [7]; the other is based on arc reversal [12, 14]. The main difference is that the standard algorithms operate with ordinary potentials, i.e., potentials that assign a real number to each configuration of their variables, while in our algorithms the conditional probabilities are ordinary potentials, but the utility is a CEP-potential, i.e., a potential that assigns a CEP to each configuration of its variables.

The algorithms presented in this paper are two adaptations of the method for performing CEA in decision trees with embedded decision nodes [1]. The main contribution of this paper is the possibility of solving IDs whose equivalent decision trees have prohibitive sizes. Using a Java implementation of the variable elimination algorithm, we have performed CEA on an ID for the mediastinal staging of non-small cell lung cancer that contains 5 decisions and 8 chance variables [9]. More recently, we have applied the same method to an ID for total knee replacement prosthesis that contains 4 decisions and 11 chance variables [8]. The equivalent decision trees, which can be obtained automatically from the influence diagrams, have thousands of leaves. Clearly, it would have been impossible to build those trees directly, and their evaluation would have been much less efficient than the evaluation of the IDs. An open line for future research is to summarize the results of these evaluations into small policy trees, using a method similar to the one proposed in [9].

Another line for future research is the adaptation of our CEA algorithms to the evaluation of decision analysis networks (DANs) [3], which present several advantages over IDs, especially in the case of asymmetric decision problems.

Acknowledgments

This work has been supported by grants TIN2006-11152 and TIN2009-09158, of the Spanish Ministry of Science and Technology, and by FONCICYT grant 85195.

References

1. M. Arias and F. J. Díez. Cost-effectiveness analysis with sequential decisions. Technical Report CISIAD-11-01, UNED, Madrid, Spain, 2011. http://www.cisiad.uned.es/techreports/cea-multidec.php.
2. C. Bielza, M. Gómez, and P. P. Shenoy. A review of representation issues and modelling challenges with influence diagrams. *Omega*, 39:227–241, 2011.
3. F. J. Díez and M. Luque. Representing decision problems with Decision Analysis Networks. Technical Report CISIAD-10-01, UNED, Madrid, Spain, 2010.
4. M. F. Drummond, M. J. Sculpher, G. W. Torrance, B. J. O'Brien, and G. L. Stoddart. *Methods for the Economic Evaluation of Health Care Programmes*. Oxford University Press, third edition, 2005.
5. M. R. Gold, J. E. Siegel, L. B. Russell, and M. C. Weinstein. *Cost-Effectiveness in Health and Medicine*. Oxford University Press, New York, 1996.
6. 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.

7. F. V. Jensen and T. D. Nielsen. *Bayesian Networks and Decision Graphs*. Springer-Verlag, New York, second edition, 2007.
8. D. León. An influence diagram for total knee arthroplasty. Master's thesis, Dept. Artificial UNED, Madrid, Spain, 2011.
9. M. Luque. *Probabilistic Graphical Models for Decision Making in Medicine.* PhD thesis, UNED, Madrid, 2009.
10. M. Luque and F. J. Díez. Variable elimination for influence diagrams with super-value nodes. *International Journal of Approximate Reasoning*, 51(6):615 – 631, 2010.
11. S. H. Nielsen, T. D. Nielsen, and F. V. Jensen. Multi-currency influence diagrams. In A. Salmerón and J. A. Gámez, editors, *Advances in Probabilistic Graphical Models*, pages 275–294. Springer, Berlin, Germany, 2007.
12. S. M. Olmsted. *On Representing and Solving Decision Problems.* PhD thesis, Dept. Engineering-Economic Systems, Stanford University, CA, 1983.
13. H. Raiffa and R. Schlaifer. *Applied Statistical Decision Theory.* MIT press, Cambridge, 1961.
14. R. D. Shachter. Evaluating influence diagrams. *Operations Research*, 34:871–882, 1986.
15. H. C. Sox, M. A. Blatt, M. C. Higgins, and K. I. Marton. *Medical Decision Making.* Butterworth-Heinemann, Woburn, MA, 1988.
16. A. A. Stinnett and J. Mullahy. Net health benefit: A new framework for the analysis of uncertainty in cost-effectiveness analysis. *Medical Decision Making*, 18:S68–S80, 1998.
17. J. A. Tatman and R. D. Shachter. Dynamic programming and influence diagrams. *IEEE Transactions on Systems, Man, and Cybernetics*, 20:365–379, 1990.
18. M. C. Weinstein and W. B. Stason. Foundations of cost-effectiveness analysis for health and medical practices. *New England Journal of Medicine*, 296:716–721, 1977.
19. M. C. Weinstein, G. Torrance, and A. McGuire. QALYs: The basics. *Value in Health*, 12(Supplement I):S5–S9, 2009.