RC25100 (R1101-001) January 18, 2011
Mathematics

IBM Research Report

Order-of-Magnitude Influence Diagrams

Radu Marinescu
IBM Research
Smarter Cities Technology Centre
Mulhuddart
Dublin 15, Ireland

Nic Wilson
Cork Constraint Computation Centre
University College Cork
Cork, Ireland

Order-of-Magnitude Influence Diagrams

Radu Marinescu
IBM Research
Smarter Cities Technology Centre
Mulhuddart, Dublin 15, Ireland
radu.marinescu@ie.ibm.com

Nic Wilson
Cork Constraint Computation Centre
University College Cork
Cork, Ireland
n.wilson@4c.ucc.ie

December 15, 2010

Abstract

Influence diagrams are a widely used framework for decision making under uncertainty. These models allow for a concise graphical representation of both probabilistic as well as utility information which in turn supports efficient graph-based algorithms for computing an optimal decision policy that maximizes the expected utility of the decision maker. In this paper we extend the framework to incorporate a qualitative rather than quantitative representation of the information based on order-of-magnitude probability and utility function. We also develop a variable elimination algorithm that generates an order-of-magnitude optimal policy. Furthermore, our model supports totally as well as partially ordered utilities. Numerical experiments on random influence diagrams analyze the quality of the order-of-magnitude policy with respect to the optimal policy derived from a corresponding regular influence diagram with exact probabilities and utility values.

1 Introduction

An influence diagram is a graphical model for decision making under uncertainty. It is composed by a directed acyclic graph where utility nodes are associated to profits and costs of actions, chance nodes represent uncertainties and dependencies in the domain and decision nodes represents actions to be taken. Given an influence diagram, a policy defines which decision to take at each node, given the information available at that moment. Each policy has a corresponding expected utility and the most common task is to find an optimal policy with maximum expected utility.

Over the past decades, several exact methods have been proposed to solve influence diagrams using local computations [1, 2, 3, 4]. These methods adapted classical *variable elimination* techniques, which compute a type of marginalization over a combination of local functions, in order to handle the multiple types of information (probabilities and utilities), marginalizations (sum and max) and combinations ($\times$ for probabilities, $+$ for utilities) involved in influence diagrams. Variable elimination based techniques are known to exploit the conditional independencies encoded by the influence diagram, however, they require time and space exponential in the *constrained treewidth* of the influence diagram.

An alternative approach for evaluating influence diagrams is based on *conditioning* (or *search*). These methods unfold the influence diagram into a *decision graph* (or *tree*) in such a way that an optimal solution graph corresponds to an optimal policy of the influence diagram. In this case, the problem of computing an optimal policy is reduced to searching for an optimal solution of the decision graph [5, 6, 7]. In contrast with variable elimination, search algorithms are not sensitive to the problem structure, use time exponential in the number of variables, but may operate in linear space.

In recent years, a number of proposals have been extended for the purpose of relieving domain experts from having to specify point probability as well as utility values. Many of these proposals offer concrete methods that allow Bayesian reasoning as well as decision making under uncertainty to commence without a commitment to complete probability distributions and utility functions. An example of this is Qualitative Probabilistic Networks [8], which allow one to reason about probabilistic influences among variables in a qualitative manner that is consistent with Bayesian reasoning. A second class of proposals attempts to relief experts from providing point probabilities by requiring more abstract and intuitive belief measures which are consistent with point probabilities. In this direction, a key proposal is the *kappa calculus* [9, 10] and its probabilistic interpretation using ϵ-semantics [11]. In the context of decision making, the Possibilistic Influence Diagrams [12] allow to model in a compact form problems of sequential decision making under uncertainty, when only ordinal data on transitions likelihood or preferences are available.

In this paper, we propose a new framework for qualitative decision making under uncertainty based on an Order of Magnitude calculus [13]. In particular, we introduce the Order of Magnitude Influence Diagram (OOMID) that extends the regular influence diagram by allowing one to work with different forms of uncertainty and other notions of utility. The graphical part of a OOMID is exactly the same as that of usual influence diagrams, however the semantics differ. Transition likelihoods are expressed as order-of-magnitude probability functions and rewards are replaced by order-of-magnitude utility functions as well. We also develop a variable elimination algorithm for computing the optimal policy that maximizes the expected utility. We consider both totally ordered as well as partially ordered utilities, thus allowing one to model the utility function in terms of both costs (negative values) and benefits (positive values). Numerical experiments on random influence diagrams analyze the quality of the order-of-magnitude policy with respect to the optimal policy derived from a corresponding regular influence diagram with exact probabilities and utility values.

2 Background

Influence diagrams extend belief networks by adding also *decision variables* and reward functional components. Formally, an influence diagram is defined by $ID = \langle X, D, P, R \rangle$, where $X = \{X_1, ..., X_n\}$ is a set of chance variables on multi-valued domains (the belief network part) and $D = \{D_1, ..., D_m\}$ is a set of decision variables (or actions). The chance variables are further divided into *observable* meaning they will be observed during the execution, or *unobservable*. The discrete domains of the decision variables denote its possible set of actions. An action at the decision variable D_i is denoted by d_i. Every chance variable X_i is associated with a conditional probability (CPT) $P_i = P(X_i|pa_i)$, where $pa_i \subseteq X \cup D - \{X_i\}$. Each decision variable D_i has a parent set $pa_{D_i} \subseteq X \cup D$, denoting the variables whose values will be known and may affect directly the decision. The reward functions $R = \{v_1, ..., v_j\}$ are defined over subsets of variables $Q = \{Q_1, ..., Q_j\}$, $Q_j \subseteq X \cup D$, called *scopes*, and the utility function is defined by $u(x) = \sum_j v_j(x_{Q_j})$.

The graph of an ID contains nodes for chance variables (drawn as circles), decision variables (drawn as rectangles) and for reward components (drawn as diamonds). The arcs in an ID can be partitioned into three disjoint sets, corresponding to the type of nodes they go into. Arcs into reward (or value) nodes represent functional dependencies by indicating the scope of the associated reward component. Arcs into chance nodes, denoted *dependency arcs*, represent probabilistic dependencies, whereas arcs into decision nodes, denoted *informational arcs*, imply information precedence; if $D_k \in D$ and there is a directed arc from $Y \in X \cup D$ to D_k, then the state of Y is known when decision D_k is made.

A *decision rule* for a decision variable D_i is a mapping:

$$\delta_i : \Omega_{pa_{D_i}} \rightarrow \Omega_{D_i}$$

where for $S \subseteq X \cup D$, Ω_S is the cross product of the individual domains of the variables in S. A *policy* is a list of decision rules $\Delta = (\delta_1, ..., \delta_m)$ consisting of one rule for each decision variable. To evaluate an influence diagram is to find the *optimal policy* that maximizes the expected utility (MEU) and to compute the optimal expected utility. Assume that x is an assignment over both chance and decision variables $x = (x_1, ..., x_n, d_1, ..., d_m)$, the MEU task is to compute:

$$E = max_{\Delta=(\delta_1,...,\delta_m)} \sum_{x_1,...,x_n} \prod P(x_i, e|x^{\Delta}_{pa_i}) \times u(x^{\Delta})$$

where x^{Δ} denotes an assignment $x = (x_1, ..., x_n, d_1, ..., d_m)$ where each d_i is determined by $\delta_i \in \Delta$ as a function of x, namely $d_i = \delta_i(x)$, and e is an instantiated subset of variables.

The set of informational arcs induces a *partial order* $\prec$ on $X \cup D$ as defined by the transitive closure of the following relation:

- $y \prec D_i$, if $(Y \rightarrow D_i)$ is a directed arc in ID ($D_i \in D$).
- $D_i \prec Y$, if $(D_i, X_1, X_2, ..., X_p, Y)$ is a directed path in ID ($Y \in X \cup D$ and $D_i \in D$).

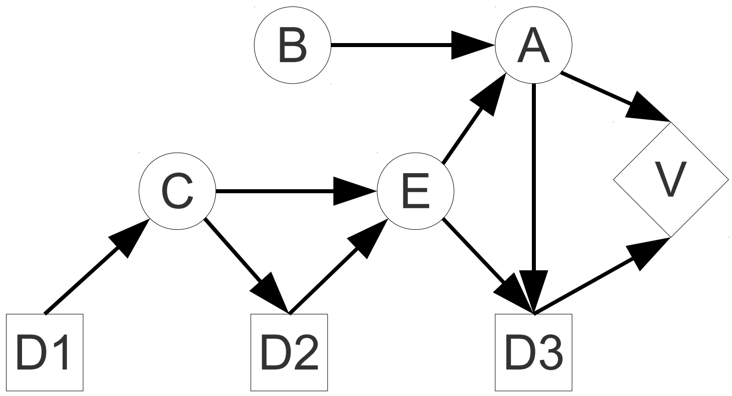

(a) Influence diagram

B	E	A	P(A|B,E)
0	0	0	0.003
0	0	1	0.997
0	1	0	0.9998
0	1	1	0.0002
1	0	0	0.992
1	0	1	0.008
1	1	0	0.99996
1	1	1	0.00004

C	D2	E	P(E|C,D2)
0	0	0	0.00007
0	0	1	0.99993
0	1	0	0.882
0	1	1	0.118
1	0	0	0.003
1	0	1	0.997
1	1	0	0.9996
1	1	1	0.0004

D1	C	P(C|D1)
0	0	0.005
0	1	0.995
1	0	0.9998
1	1	0.0002

B	P(B)
0	0.99
1	0.01

A	D3	V(A,D3)
0	0	-3000
0	1	-50
1	0	0
1	1	80

(b) Conditional probability and utility functions

Figure 1: An influence diagram with elimination order D1, C, D2, E, A, D3, B.

- $D_i \prec A$, if $A \nprec D_j$ for all $D_j \in D$ ($A \in X$ and $D_i \in D$).
- $D_i \prec A$, if $A \nprec D_i$ and $\exists D_j \in D$ s.t. $D_i \prec D_j$ and $A \prec D_j$ ($A \in X$ and $D_i \in D$).

Example 1 *Figure 1 shows an influence diagram with 3 decision variables (D1, D2, D3), 4 chance variables (A,B,C,E) and 1 value node (V), respectively. The partial order induced by the influence diagram is:* $D1 \prec C \prec D2 \prec \{A, E\} \prec D3 \prec B$. *The optimal policy of the diagram has maximum expected utility 78.19754 and is represented by the decision rules:* $\delta(D1) :\rightarrow \Omega_{D1}$, $\delta(D2) : \{C\} \rightarrow \Omega_{D2}$ *and* $\delta(D3) : \{A\} \rightarrow \Omega_{D3}$, *respectively, where:*

$$\delta(D1) = 0$$

$$\delta(D2) = \begin{cases} 1 & \textit{if } C = 0; \\ 1 & \textit{if } C = 1; \end{cases}$$

$$\delta(D3) = \begin{cases} 1 & \textit{if } A = 0; \\ 1 & \textit{if } A = 1; \end{cases}$$

3 Order-of-Magnitude Calculus

In this section, we introduce the Order-of-Magnitude calculus as a representation framework for imprecise probabilities as well as imprecise utilities.

3.1 Definition of Order of Magnitude Calculus

Let $\mathcal{O} = \{\langle \sigma, n\rangle \,:\, n \in \mathbb{Z}, \sigma \in \{+, -, \pm\}\} \cup \{\langle 0, \infty\rangle\}$, where $\mathbb{Z}$ is the set of integers. The element $\langle \pm, \infty\rangle$ will sometimes be written as 0, element $\langle +, 0\rangle$ as 1, and element $\langle -, 0\rangle$ as -1.

If $a = \langle \sigma, n\rangle$ then define $\sigma(a)$ (the sign of a) to be σ, and $\hat{a}$ to be n. We also define $\mathcal{O}_{\pm} = \{\langle \pm, n\rangle \,:\, n \in \mathbb{Z} \cup \{\infty\}\}$, and $\mathcal{O}_{+} = \{\langle +, n\rangle \,:\, n \in \mathbb{Z}\}$.

Multiplication:

For $\langle \sigma, m\rangle, \langle \sigma', n\rangle \in \mathcal{O}$, let $\langle \sigma, m\rangle \times \langle \sigma', n\rangle = \langle \sigma \otimes \sigma', m + n\rangle$, where $\infty + m = m + \infty = \infty$ for $m \in \mathbb{Z} \cup \{\infty\}$, and $\otimes$ is the natural multiplication of signs: it is the commutative operation on $\{+, -, \pm\}$ such that $+ \otimes - = -$, $+ \otimes + = - \otimes - = +$, and for any $\sigma \in \{+, -, \pm\}$, $\sigma \otimes \pm = \pm$. As usual, $a \times b$ will be sometimes abbreviated to ab. This multiplication is associative and commutative, and $(\mathcal{O} \setminus \mathcal{O}_{\pm}, \times)$ is an abelian group. Also $-1 \times -1 = 1$ and for any $a \in \mathcal{O}$, $a \times 0 = 0$ and $a \times 1 = a$.

For $b \in \mathcal{O} \setminus \mathcal{O}_{\pm}$ define b^{-1} to be the multiplicative inverse of b, and for $a \in \mathcal{O}$ let $a/b = a \times b^{-1}$. $\langle \sigma, m\rangle^{-1} = \langle \sigma, -m\rangle$ for $\sigma \in \{+, -\}$.

Addition:

For $\langle \sigma, m\rangle, \langle \sigma', n\rangle \in \mathcal{O}$, let

$$\langle \sigma, m\rangle + \langle \sigma', n\rangle = \begin{cases} \langle \sigma, m\rangle & \text{if } m < n; \\ \langle \sigma', n\rangle & \text{if } m > n; \\ \langle \sigma \oplus \sigma', m\rangle & \text{if } m = n \end{cases}$$

where $+ \oplus + = +$, $- \oplus - = -$, and otherwise, $\sigma \oplus \sigma' = \pm$.

Addition is associative and commutative, and $a + 0 = a$ for any $a \in \mathcal{O}$. We have distributivity: for $a, b, c \in \mathcal{O}$, $(a + b)c = ac + bc$.

For $a, b \in \mathcal{O}$ let $-b = -1 \times b$, and $a - b = a + (-b)$. We have $-\langle \sigma, m\rangle = \langle -\sigma, m\rangle$ where, as one would expect, $-(+) = -$, $\quad -(-) = +$ and $-(\pm) = \pm$.

Ordering

We use a slightly stronger ordering than that defined in [13].

Let a and b be some elements of $\mathcal{O}$. Write a as $\langle\sigma, m\rangle$ and b as $\langle\tau, n\rangle$. Then we define binary relation $\preceq$ on $\mathcal{O}$ by $a \succeq b$ if and only if either:

- $\sigma = +$ and $\tau = +$ and $m \leq n$; or
- $\sigma = +$ and $\tau = \pm$ and $m \leq n$; or
- $\sigma = +$ and $\tau = -$; or
- $\sigma = \pm$ and $\tau = -$ and $m \geq n$; or
- $\sigma = -$ and $\tau = -$ and $m \geq n$.

As usual we write $a \succ b$ if and only if $a \succeq b$ and it is not the case that $b \succeq a$. We write $a \preceq b$ if and only if $b \succeq a$, and $a \prec b$ if and only if $b \succ a$.

We sum up some basic properties of $\succeq$:

Proposition 1 *Let a, b, c be arbitrary elements of $\mathcal{O}$.*

- *$\succeq$ is a partial order on $\mathcal{O}$.*
- *$a \succeq b$ if and only if $-b \succeq -a$.*
- *$a \succ 0$ if and only if $a \in \mathcal{O}_+$.*
- *If $a \succeq b$ then $a + c \succeq b + c$.*
- *If $a \succeq b$ and $c \in \mathcal{O}_+$ then $a \times c \succeq b \times c$.*

3.2 Lower Simplified Order of Magnitude Calculus

We define the Lower Simplified Order of Magnitude Calculus as follows.

Define the set L to consist of pairs $\langle\sigma, n\rangle$ where $\sigma \in \{+, -\}$ and n is an integer.

Multiplication:

$(\sigma, m) \times (\sigma', n) = (\sigma \boxtimes \sigma', m + n)$, where $\boxtimes$ is the natural multiplication of signs: it is the commutative operation on $\{+, -\}$ such that $+ \boxtimes - = -$, $+ \boxtimes + = - \boxtimes - = +$. This multiplication is associative and commutative, and has inverses: $(\sigma, m)^{-1} = (\sigma, -m)$ for $\sigma \in \{+, -\}$.

Addition: Let

$$(\sigma, m) + (\sigma', n) = \begin{cases} (\sigma, m) & \text{if } m < n; \\ (\sigma', n) & \text{if } m > n; \\ (\sigma \boxplus \sigma', m) & \text{if } m = n \end{cases}$$

where $+ \boxplus + = +$, and otherwise, $\sigma \boxplus \sigma' = -$.

3.3 Upper Simplified Order of Magnitude Calculus

Given a set L consisting of pairs $\langle \sigma, n, \rangle$, where $\sigma \in \{+, -\}$ and n is an integer, we define the Upper Simplified Order of Magnitude Calculus (SUOOM) in a similar way, as follows.

Multiplication:
$(\sigma, m) \times (\sigma', n) = (\sigma \boxtimes \sigma', m + n)$, where $\boxtimes$ is the natural multiplication of signs: it is the commutative operation on $\{+, -\}$ such that $+ \boxtimes - = -$, $+ \boxtimes + = - \boxtimes - = +$. This multiplication is associative and commutative, and has inverses: $(\sigma, m)^{-1} = (\sigma, -m)$ for $\sigma \in \{+, -\}$.

Addition: Let

$$(\sigma, m) + (\sigma', n) = \begin{cases} (\sigma, m) & \text{if } m < n; \\ (\sigma', n) & \text{if } m > n; \\ (\sigma \boxplus \sigma', m) & \text{if } m = n \end{cases}$$

where $+ \boxplus + = +$, and otherwise, $\sigma \boxplus \sigma' = +$.

3.4 Ordering for the Simplified Calculus

For the simplified (upper and lower) order of magnitude calculus we can define a total order as follows. Let L be a set of elements of the form $\langle \sigma, n \rangle$, where n is an integer.

Ordering

$\langle \sigma, m \rangle \succeq \langle \tau, n \rangle$ if and only if either

- $\sigma = +$ and $\tau = -$; or
- $\sigma = +$ and $\tau = +$, and $m \leq n$; or
- $\sigma = -$ and $\tau = -$, and $m \geq n$.

Lemma 1 *$\succeq$ respects addition $+$.*

For the totally ordered case we have a total order $\succeq$ on the set U of utilities. The max-marginalization operation on utility-value pairs uses the operation $\max_{\succeq}$, returns the maximum (with respect to $\succeq$) of a pair of utility values.

In order for combination of pairs to distribute over the max-marginalization we need that addition of utilities $+$ distributes over $\max_{\succeq}$. We say that $\succeq$ respects $+$ if the following (monotonicity) property holds:

$\forall a, b, c \in U$, if $a \succeq b$ then $a + c \succeq b + c$;

Theorem 1 *Given total order $\succeq$ on U which respects $+$ (addition of utilities), the operation $\max_{\succeq}$ is commutative and associative, and $+$ distributes over $\max_{\succeq}$.*

4 Max Marginalization over Sets of Partially Ordered Order-of-Magnitude Utilities

4.1 Ordering on utilities

We assume a partial ordering $\succeq$ on U which satisfies the following monotonicity properties, i.e., $\succeq$ respects $+$ and $\times$:

- for all $a, b, c \in U$, if $a \succeq b$ then $a + c \succeq b + c$;
- for all $a, b \in U$ and for all $q \in Q$, if $a \succeq b$ then $q \times a \succeq q \times b$.

If $a \succeq b$ then we say that a dominates b. We write $\succ$ for the strict part of $\succeq$, so that $a \succ b$ if and only if $a \succeq b$ and $a \neq b$. For $A \subseteq U$, define $\max_{\succeq}(A)$, *the maximal elements of* A, to consist of all $a \in A$ such that there does not exist $b \in A$ with $b \succ a$. Hence $\max_{\succeq}(A)$ is the set of undominated element of A.

In many cases, every element of a set A is dominated by some maximal element; in particular this holds if A is finite. This can allow a set of utilities to be summarized by its maximal elements. However, this is not universally true, since we may have infinite chains $a_1 \prec a_2 \prec a_3 \prec \cdots$ which have no upper bound in A. Consider, for example, the open interval $(0, 1)$ of the real numbers, which has no maximal elements.

Definition 1 *Given partial ordering $\succeq$ on set A, we say that A satisfies property MAX if for all $a \in A$ there exists some $b \in \max_{\succeq}(A)$ with $b \succeq a$.*

For any finite set A and any partial order $\succeq$ on A, then A satisfies MAX. For $A \subseteq U$ we define subset $\mathcal{R}_{\succeq}(A)$ of A to consist of all elements of A which are not strictly dominated by some maximal element of A, that is:

$$\mathcal{R}_{\succeq}(A) = \{a \in A \,:\, \nexists b \in \max_{\succeq}(A) \quad \text{such that} \quad b > a\}.$$

Clearly, we always have $\max_{\succeq}(A) \subseteq \mathcal{R}_{\succeq}(A)$. If A is such that every element of A is dominated by some maximal element of A (in particular, this is the case if A is finite), then $\mathcal{R}_{\succeq}(A) = \max_{\succeq}(A)$.

Lemma 2 *Let $\succeq$ be a partial order on set A.*

(i) $\max_{\succeq}(A) \subseteq \mathcal{R}_{\succeq}(A)$

(ii) If A satisfies MAX then $\max_{\succeq}(A) = \mathcal{R}_{\succeq}(A)$.

Lemma 3 *Let A and B be subsets of U.*

$\max_{\succeq}(\max_{\succeq}(A) \cup B)) = \max_{\succeq}(A \cup B)$.

Lemma 4 *Let A and B be subsets of U.*

$\max_{\succeq}(\mathcal{R}_{\succeq}(A) \cup B)) = \max_{\succeq}(A \cup B)$.

Ordering on sets of utilities

For $A, B \subseteq U$ we say that $A \succcurlyeq B$ if every element of B is dominated by some element of A (so that A contains as least as large elements as B), i.e., if for all $b \in B$ there exists $a \in A$ with $a \succeq b$. Moreover, $\succcurlyeq$ is a reflexive and transitive relation. We define equivalence relation $\approx$ by $A \approx B$ if and only if $A \succcurlyeq B$ and $B \succcurlyeq A$.

If A is such that every element of A is dominated by some maximal element of A (for example, if A is finite), then $A \approx B$ if and only if $\max_{\succeq}(A) = \max_{\succeq}(B)$.

Lemma 5 *Let $\succeq$ be a partial order on set A, which satisfies MAX. Then*

$A \approx B$ *if and only if* $\max_{\succeq}(A) = \max_{\succeq}(B)$.

Lemma 6 *Let A, B and C be subsets of U, and let q be an element of Q. Suppose that $A \approx B$. Then*

(i) $q \times A \approx q \times B$*;*

(ii) $A + C \approx B + C$*;*

4.2 The Equivalence Relation $\equiv$ Between Utility Sets

Throughout this section we assume a set U of order-of-magnitude utility values and that $\succeq$ is a partial order on U which respects $+$ and respects $\times_{QU}$. We define relation $\equiv$ on subsets of U by: $A \equiv B$ if and only if $\mathcal{C}(A) \approx \mathcal{C}(B)$, where $\mathcal{C}(X)$ is the convex closure of the subset $X \subseteq U$.

The definition immediately implies that $\equiv$ is an equivalence relation, i.e., it is reflexive, symmetric and transitive, since $\approx$ is an equivalence relation. Thus two sets of utility values are considered equivalent if, for every convex combination of elements of one, there is a convex combination of elements of the other which is at least as good (with respect to the partial order $\succeq$ on U).

Proposition 2 *For any subset A of U, $A \equiv \mathcal{R}_{\succeq}(A)$ and $A \equiv \mathcal{C}(A)$. If A satisfies MAX (in particular, if A is finite) then $A \equiv \max_{\succeq}(A)$.*

Proposition 3 *If A and $\mathcal{C}(A)$ satisfy property MAX then* $\mathcal{R}(\mathcal{C}(A)) = \max_{\succeq}(\mathcal{C}(A)) = \max_{\succeq}(\mathcal{C}(\max_{\succeq}(A))$.

Combining this result with Lemma 5 gives:

Proposition 4 *If $A, B, \mathcal{C}(A)$ and $\mathcal{C}(B)$ satisfy property MAX then $A \equiv B$ if and only if* $\max_{\succeq}(\mathcal{C}(A)) = \max_{\succeq}(\mathcal{C}(B))$.

Proposition 5 *Let A, B and C be subsets of U, and let q be an element of Q. Suppose that $A \equiv B$. Then*

(i) $q \times A \equiv q \times B$*;*

(ii) $A + C \equiv B + C$*;*

(iii) $A \cup C \equiv B \cup C$.

4.3 Generating Small Equivalent Sets

The key result in this section is Proposition 7, that implies that in the (partially ordered) Order of Magnitude computation (OOM), one needs only to work with sets of values which have either one or two elements. The result refers to the equivalence relation $\equiv$ defined by $A \equiv B$ if and only if $\mathcal{C}(A) \approx \mathcal{C}(B)$.

Consider two elements $\langle \sigma, m \rangle$ and $\langle \tau, n \rangle$ in $\mathcal{O}$, where we can assume, without loss of generality, that $m \leq n$. Any convex combination of these two elements is of the form $\langle \theta, l \rangle$ where $l \in [m, n]$ and

if $l < n$ then $\theta = \sigma$;

if $l = n$ then $\theta = \sigma \boxplus \tau$ or $\theta = \tau$.

This means that convex sets are of a relatively simple form. In particular, it implies that the convex combination of a finite number of non-zero elements is finite (since every element a in the convex combination has $\hat{a}$ restricted to be within a finite range), and so, in particular, satisfies the important MAX property (see Definition 1 in Section 4.1). In fact, even if we allow the zero element $\langle \pm, \infty \rangle$, $\mathcal{C}(A)$ still satisfies MAX:[1]

Lemma 7 *Let A be a finite subset of $\mathcal{O}$. Then A and $\mathcal{C}(A)$ both satisfy the MAX property*

Lemma 7, Proposition 3, and Proposition 4 imply the following result, which gives a simpler definition of equivalence.

Proposition 6 *Let A and B be finite subsets of $\mathcal{O}$. Then $A \equiv B$ if and only if $\max_{\succeq}(\mathcal{C}(A)) = \max_{\succeq}(\mathcal{C}(B))$. Also, $\max_{\succeq}(\mathcal{C}(A)) = \max_{\succeq}(\mathcal{C}(\max_{\succeq}(A))$, and $\max_{\succeq}(\mathcal{C}(B)) = \max_{\succeq}(\mathcal{C}(\max_{\succeq}(B))$.*

In fact, it turns out that $\max_{\succeq}(\mathcal{C}(\max_{\succeq}(A)) = \mathcal{C}(\max_{\succeq}(A))$, and similarly for B. To prove Proposition 7 below we use Lemmas 8 and 9.

Lemma 8 *Let A be any subset of $\mathcal{O}$ with $\max_{\succeq}(A) = A$. Then there is at most one element $a \in A$ with $\sigma(a) \neq \pm$, i.e.., with $\sigma(a) = +$ or $-$. Furthermore, If $a \in A$ is such that $\sigma(a) \neq \pm$ then for all other elements b of A, $\hat{b} < \hat{a}$.*

Lemma 9 *Let A be any finite subset of $\mathcal{O}$ with $\max_{\succeq}(A) = A$. Then either $|A| = 1$ or there exists some $m, n \in \mathbb{Z}$ with $m < n$ and $\sigma \in \{+, -, \pm\}$ such that $\mathcal{C}(A) = \mathcal{C}(\{\langle \pm, m \rangle, \langle \sigma, n \rangle\})$.*

Proposition 7 *Let A be any finite subset of $\mathcal{O}$. Then either $A \equiv \{a\}$ for some $a \in \mathcal{O}$, or there exists some $m, n \in \mathbb{Z}$ with $m < n$ and $\sigma \in \{+, -, \pm\}$ such that $A \equiv \{\langle \pm, m \rangle, \langle \sigma, n \rangle\}$.*

This implies that, when computing with pairs (q, A), in order to perform variable elimination for OOM-based influence diagrams, we always replace set A by a set A' which has either one or two elements, such that $A' \equiv A$. This affects the complexity of the procedure, which is related to the size of sets A that are used in the computation.

[1] Even though $\mathcal{C}(A)$ is not necessarily finite: consider e.g., $\mathcal{C}(\{\langle \pm, \infty \rangle, \langle -, m \rangle\})$ which includes all elements $\langle -, n \rangle$ with $n \geq m$.

5 Order-of-Magnitude Influence Diagrams

An order-of-magnitude influence diagram (OOM-ID) is a formalism for decision making using imprecise probabilistic as well as imprecise utility information by considering an order-of-magnitude approximation of the probabilities and utilities, respectively. Specifically, the graphical structure of an *order-of-magnitude influence diagram* is identical to that of a regular influence diagram, in terms of chance nodes, decision nodes, value (or utility) nodes, as well as the dependency relationships among them. However, the entries of the conditional probability tables as well as of the utility functions are represented as order-of-magnitude values, namely elements of the form $\langle\sigma, n\rangle$, where n is an integer and $\sigma \in \{+, -, \pm\}$ is the sign, respectively.

One can always approximate a regular influence diagram by an order-of-magnitude influence diagram using the following conversion procedure of the exact probability and utility values. This conversion which only maintains the order of magnitude of the respective probability or utility value, is a way of capturing some degree of imprecision and was also used previously by [14] in the context of Bayesian networks.

Given a point probability value $p \in [0, 1]$ and $0 < \epsilon \leq 1$, the corresponding order-of-magnitude approximation is $(+, n)$ where

$$n = \begin{cases} \lfloor \log_\epsilon p \rfloor & \text{if } 0 < p \leq 1; \\ \infty & \text{if } p = 0; \end{cases}$$

Similarly, given a point utility value $u \in \mathbb{R}$ and $0 < \epsilon \leq 1$, the corresponding order-of-magnitude approximation is: (σ, n) where

$$n = \begin{cases} \lceil \log_\epsilon |u| \rceil & \text{if } |u| \geq 1; \\ \lfloor log_\epsilon |u| \rfloor & \text{if } 0 < |u| < 1; \\ \infty & \text{if } u = 0 \end{cases}$$

and

$$\sigma = \begin{cases} + & \text{if } u \geq 0; \\ - & \text{if } u < 0; \end{cases}$$

Example 2 *For illustration, Figure 2 shows the order-of-magnitude approximation of conditional probability tables and utility (reward) components of the influence diagram of Figure 1 for* $\epsilon = 0.1$. *For example, the OOM approximation of the point probability* $P(A = 0|B = 0, E = 0) = 0.002$ *is computed as* $(+, \lfloor \log_{0.1} 0.002 \rfloor) = (+, \lfloor \frac{\log 0.002}{\log 0.1} \rfloor) = (+, \lfloor 2.698 \rfloor) = (0, 2)$. *Similarly, the OOM approximation of the point utility* $V(A = 0, D3 = 0) = -3000$ *is computed as* $(-, \lceil \log_{0.1} 3000 \rceil) = (-, \lceil \frac{\log 3000}{\log 0.1} \rceil) = (-, -3)$.

Example 3 *For* $\epsilon = 0.1$, *the OOM approximation of the influence diagram in Example 1 has the following optimal policy with maximum expected utility* $(+, -1)$, *which is identical to the optimal policy of the original influence diagram:*

$$\delta(D1) = 0$$

B	E	A	P(A|B,E)
0	0	0	(+, 2)
0	0	1	(+, 0)
0	1	0	(+, 0)
0	1	1	(+, 3)
1	0	0	(+, 0)
1	0	1	(+, 2)
1	1	0	(+, 0)
1	1	1	(+, 4)

C	D2	E	P(E|C,D2)
0	0	0	(+, 4)
0	0	1	(+, 0)
0	1	0	(+, 0)
0	1	1	(+, 0)
1	0	0	(+, 2)
1	0	1	(+, 0)
1	1	0	(+, 0)
1	1	1	(+, 3)

D1	C	P(C|D1)
0	0	(+, 2)
0	1	(+, 0)
1	0	(+, 0)
1	1	(+, 3)

B	P(B)
0	(+, 0)
1	(+, 1)

A	D3	V(A,D3)
0	0	(-, -3)
0	1	(-, -1)
1	0	(+, ∞)
1	1	(+, -1)

Figure 2: Order-of-magnitude approximation of the influence diagram from Figure 1

$$\delta(D2) = \begin{cases} 1 & \textit{if } C = 0; \\ 1 & \textit{if } C = 1; \end{cases}$$

$$\delta(D3) = \begin{cases} 1 & \textit{if } A = 0; \\ 1 & \textit{if } A = 1; \end{cases}$$

For $\epsilon = 0.001$, the OOM approximation of the influence diagram in Example 1 has the following optimal policy with maximum expected utility $(\pm, 0)$:

$$\delta(D1) = \{0, 1\}$$

$$\delta(D2) = \begin{cases} 1 & \textit{if } C = 0; \\ \{0, 1\} & \textit{if } C = 1; \end{cases}$$

$$\delta(D3) = \begin{cases} 1 & \textit{if } A = 0; \\ 1 & \textit{if } A = 1; \end{cases}$$

Note that in the latter case, both decision alternatives are allowed by $\delta(D1)$ and by $\delta(D2)$ for $C = 1$, respectively. Namely, the more imprecise one is about the conditional probabilities and/or utility values (corresponding to smaller ϵ values), the number of undecided decision rules in the optimal policy of the OOM approximation increases.

In this case, the optimal strategy of the order-of-magnitude influence diagram is represented by a policy set $\Delta_s = \{\Delta_1, \Delta_2, \Delta_3, \Delta_4\}$ *containing the following policies:*

The policy $\Delta_1 = \{\delta_1(D1), \delta_1(D2), \delta_1(D3)\}$ is:

$$\delta_1(D1) = 0$$

$$\delta_1(D2) = \begin{cases} 1 & \textit{if } C = 0; \\ 0 & \textit{if } C = 1; \end{cases}$$

$$\delta_1(D3) = \begin{cases} 1 & \textit{if } A = 0; \\ 1 & \textit{if } A = 1; \end{cases}$$

The policy $\Delta_2 = \{\delta_2(D1), \delta_2(D2), \delta_2(D3)\}$ *is:*

$$\delta_2(D1) = 0$$

$$\delta_2(D2) = \begin{cases} 1 & \textit{if } C = 0; \\ 1 & \textit{if } C = 1; \end{cases}$$

$$\delta_2(D3) = \begin{cases} 1 & \textit{if } A = 0; \\ 1 & \textit{if } A = 1; \end{cases}$$

The policy $\Delta_3 = \{\delta_3(D1), \delta_3(D2), \delta_3(D3)\}$ *is:*

$$\delta_3(D1) = 1$$

$$\delta_3(D2) = \begin{cases} 1 & \textit{if } C = 0; \\ 0 & \textit{if } C = 1; \end{cases}$$

$$\delta_3(D3) = \begin{cases} 1 & \textit{if } A = 0; \\ 1 & \textit{if } A = 1; \end{cases}$$

The policy $\Delta_4 = \{\delta_4(D1), \delta_4(D2), \delta_4(D3)\}$ *is:*

$$\delta_4(D1) = 1$$

$$\delta_4(D2) = \begin{cases} 1 & \textit{if } C = 0; \\ 1 & \textit{if } C = 1; \end{cases}$$

$$\delta_4(D3) = \begin{cases} 1 & \textit{if } A = 0; \\ 1 & \textit{if } A = 1; \end{cases}$$

Notice that in the latter case, namely $\epsilon = 0.001$*, only* Δ_2 *is equal to the optimal policy of the original influence diagram. When we evaluated the policy set in the original influence diagram we obtained the following expected utilities:*

policy	*expected utility*
Δ_1	*-49.02544*
Δ_2	**78.19754**
Δ_3	*63.16252*
Δ_4	*63.18809*

However, all four decision rules are optimal in the OOM influence diagram since they are all of the same order of magnitude for $\epsilon = 0.001$*, namely* $(\pm, 0)$*.*

Algorithm 1: Variable Elimination: ELIM-OOM-ID

Data: Order-of-magnitude influence diagram $\langle X, D, P, R \rangle$, a legal elimination ordering of the variables o.

Result: The optimal policy $\Delta = \{\delta_1, ..., \delta_l\}$ that maximizes the expected utility.

`/* initialize: partition functions into buckets */`

create a set $buckets$ of size n

for *p = n downto 1* **do**

  let X_p be the p^{th} variable in ordering o and associate $buckets[p]$ to X_p

  let Λ_p be the set of functions in P that contain X_p in their scope

  if X_p *is a chance variable* **then**

    let Θ_p be the set of functions in R that contain X_p in their scope

  else if X_p *is a decision variable* **then**

    let Θ_p be the set of remaining functions in R

  place Λ_p into $buckets[p]$ and update $P \leftarrow P - \Lambda_p$

  place Θ_p into $buckets[p]$ and update $R \leftarrow R - \Theta_p$

`/* top-down phase: eliminate variables */`

for $p = n$ *downto 1* **do**

  let $\Lambda_p = \{\lambda_1, ..., \lambda_j\}$ be probabilistic components in $buckets[p]$

  let $\Theta_p = \{\theta_1, ..., \theta_k\}$ be utility components in $buckets[p]$

  if X_p *is a chance variable* **then**

    $\lambda_p \leftarrow \boxplus_{X_p} \boxtimes_{i=1}^{j} \lambda_i$

    $\theta_p \leftarrow \boxplus_{X_p}((\boxtimes_{i=1}^{j}\lambda_i) \boxtimes (\boxplus_{j=1}^{k}\theta_j))$

    $\theta_p \leftarrow \theta_p \boxtimes \lambda_p^{-1}$

  else if X_p *is a decision variable* **then**

    if $\Lambda_p = \emptyset$ **then**

      $\theta_p \leftarrow \vee_{X_p} \boxplus_{j=1}^{k} \theta_j$

    else

      $\lambda_p \leftarrow \vee_{X_p}((\boxtimes_{i=1}^{j}\lambda_i) \boxtimes (\boxplus_{j=1}^{k}\theta_j))$

  place λ_p in the bucket of the largest-index variable in its scope

  place θ_p in the closest chance bucket of a variable in its scope or in the closest decision bucket

`/* bottom-up phase: compute optimal policy */`

$\Delta \leftarrow \emptyset$

for $p = 1$ *to* n **do**

  if X_p *is a decision variable* **then**

    let $\Lambda_p = \{\lambda_1, ..., \lambda_j\}$ be probabilistic components in $buckets[p]$

    let $\Theta_p = \{\theta_1, ..., \theta_k\}$ be utility components in $buckets[p]$

    $\delta_p \leftarrow argmax_{X_p}((\boxtimes_{i=1}^{j}\lambda_i) \boxtimes (\boxplus_{j=1}^{k}\theta_j))$

    $\Delta \leftarrow \Delta \cup \delta_p$

return Δ

6 Variable Elimination for Order-of-Magnitude Influence Diagrams

A variable elimination procedure for computing the optimal policy of an order-of-magnitude influence diagram is described by Algorithm 1. Given a legal elimination ordering of the variables, the algorithm constructs a bucket structure called $buckets$ where each bucket is associated with a single variable. The input probability and utility functions are then partitioned into the buckets as follows. Each probability function is placed in the bucket of its argument that appears latest in the ordering. A utility function is placed in the bucket of its highest chance variable in its scope or in the bucket of the highest decision variable in the ordering (lines 1–10). The algorithm has two phases.

During the first, top-down phase, it processes each bucket, from the last variable to the first. Each bucket containing utility components $\Theta_p = \{\theta_1, ..., \theta_k\}$ and probability components $\Lambda_p = \{\lambda_1, ..., \lambda_j\}$, respectively, is processed by a variable elimination procedure that computes new probability and utility components which are placed in lower buckets (lines 11–24). The algorithm generates the λ_p of a bucket by combining all probability components and eliminating the bucket variable. The θ_p of a chance variable X_p is computed as the average utility of the bucket, normalized by the bucket's compiled λ_p. For a decision variable we compute a λ_p component by $\vee$-maximization, and simplify when no probabilistic components appear in the decision bucket. We note therefore that processing a decision variable does not in general allow exploiting a decomposition in the utility components. The procedure uses the following combination and elimination operators:

- $\boxtimes$-combination: $\lambda_1 \boxtimes \lambda_2 \equiv \lambda_1 \times \lambda_2$
- $\boxplus$-combination: $\theta_1 \boxplus \theta_2 \equiv \theta_1 + \theta_2$
- $\boxtimes$-combination: $\lambda \boxtimes \theta \equiv \lambda \times \theta$
- $\boxplus$-elimination: $\boxtimes_X \equiv \sum_X$
- $\vee$-elimination: this is the closure defined in the previous sections. Note the the $\vee$-elimination becomes the regular max when we consider totally ordered order-of-magnitude values (i.e., SLOOM/SUOOM simplified calculus), and is the max marginalization over partially ordered sets of order-of-magnitude values when we consider OOM calculus with both positive and negative utility values.

In the second, bottom-up phase, the algorithm computes the optimal policy or the set of decision rules for each decision variable. The decision buckets are processed in reversed order, from the first variable to the last (thus taking into account the temporal order of the decision induced by the influence diagram). Each decision rule is computed by the argument of the $\vee$-elimination operator applied over the combination of the probability and utility components in the respective bucket (including original as well as intermediate functions). We note that the scope of the decision rule (also called its *domain*) is the union of the scopes of all functions in the that bucket minus the bucket variable.

7 Empirical Evaluation

In this section, we evaluate empirically the quality of the decision policies produced by the order-of-magnitude approximation of influence diagrams.

7.1 Random problem generator

We experimented with random problems using the following parametric model. A random influence diagram class is defined by $\langle n_c, n_d, u, k_c, k_d, p, r, a \rangle$ where n_c is the number of chance variables, n_d is the number of decision variables (usually, $n_d \ll n_c$), u is the number of utility (or reward) components, k_c is the domain size of the chance variables, k_d is the domain size of the decision variables, p is the number of parents for each of the chance or decision variables, r is the number of root variables (without any parents) and a is the arity of the utility components.

The structure of the influence diagram is created by randomly picking $n_c + n_d - r$ variables out of $n_c + n_d$ and, for each, randomly selecting p parents from their preceding variables, relative to some ordering. We also ensure that the decision variables are connected by a directed path in the resulting directed acyclic graph. Then, u utility nodes are added to the graph, each one having a parents selected randomly from the set of chance and decision nodes. The fraction of chance nodes that are assigned extreme CPTs is a parameter, called the *extreme ratio* and denoted by e, respectively. The CPTs of these nodes were filled with numbers between $(10^{-5}, 10^{-4})$ ensuring that they are properly normalized; in the remaining chance nodes, the CPTs were randomly filled using a uniform distribution between 0 and 1. The table of each utility component was filled with integers v of the form:

$$v = \begin{cases} 10^i & \text{if } \ i > 0; \\ -10^i & \text{if } \ i < 0; \\ 0 & \text{if } \ i = 0; \end{cases}$$

where i is an integer uniformly distributed at random in the interval $[a, b]$, with $a, b \in \mathbb{Z}$.

7.2 Measures of performance

In order to measure the quality of the optimal policies computed in the order-of-magnitude influence diagrams we consider the following scenario.

Let $\mathcal{I}$ be a regular influence diagram and let $\mathcal{I}_\epsilon$ be its order-of-magnitude approximation for some value of ϵ obtained using the conversion procedure outlined before. Let $\Delta^*(\mathcal{I})$ be the optimal policy of $\mathcal{I}$ and let $E^*(\mathcal{I})$ be its maximum expected utility, respectively. Also, let $\Delta^*(\mathcal{I}_\epsilon) = \{\Delta^*_1(\mathcal{I}_\epsilon), \Delta^*_2(\mathcal{I}_\epsilon), ..., \Delta^*_m(\mathcal{I}_\epsilon)\}$ be the policy set of $\mathcal{I}_\epsilon$. We then draw a subset of policies $\Delta^*_s(\mathcal{I}_\epsilon) = \{\Delta^*_{i_1}(\mathcal{I}_\epsilon), ..., \Delta^*_{i_n}(\mathcal{I}_\epsilon)\}$ uniformly at random out of $\Delta^*(\mathcal{I}_\epsilon)$ and for each policy $\Delta^*_{i_j}(\mathcal{I}_\epsilon) \in D^*_s(\mathcal{I}_\epsilon)$ we compute its corresponding expected utility $E^*_{i_j}(\mathcal{I}_\epsilon)$ in the original influence diagram $\mathcal{I}$. Let $\mathcal{E}_s = \{E^*_{i_1}(\mathcal{I}_\epsilon), ..., E^*_{i_n}(\mathcal{I}_\epsilon)\}$ and we denote by $E^*_{avg}(\mathcal{I}_\epsilon)$ be the average expected utility over the sample set $\mathcal{E}_s$.

We define next the average relative error (should get a better word) as follows:

$$RE^{OOM}_{avg} = \frac{|E^*(\mathcal{I}) - E^*_{avg}(\mathcal{I}_\epsilon)|}{|E^*(\mathcal{I})|} \cdot 100$$

Similarly, we can define the minimum, maximum and median relative errors:

$$RE^{OOM}_{min} = \frac{|E^*(\mathcal{I}) - E^*_{min}(\mathcal{I}_\epsilon)|}{|E^*(\mathcal{I})|} \cdot 100$$

$$RE^{OOM}_{max} = \frac{|E^*(\mathcal{I}) - E^*_{max}(\mathcal{I}_\epsilon)|}{|E^*(\mathcal{I})|} \cdot 100$$

$$RE^{OOM}_{med} = \frac{|E^*(\mathcal{I}) - E^*_{med}(\mathcal{I}_\epsilon)|}{|E^*(\mathcal{I})|} \cdot 100$$

where $E^*_{min}(\mathcal{I}_\epsilon)$, $E^*_{max}(\mathcal{I}_\epsilon)$ and $E^*_{med}(\mathcal{I}_\epsilon)$ are the minimum, maximum and the median expected utility over the sample set $\mathcal{E}_s$, respectively.

In the following subsection we will use RE^{OOM}_{avg} (resp. RE^{OOM}_{min}, RE^{OOM}_{max} and RE^{OOM}_{med}) to characterize the quality of the order-of-magnitude decision policies with respect to the optimal decision policy of the corresponding regular influence diagram.

7.3 Results

7.3.1 Influence diagrams with positive utilities

We generated a set of random influence diagrams using the random model generator with parameters $\langle n_c, n_d = 5, u = 1, p = 2, r = 2, a = 5, e\rangle$ while varying the number of chance variables n_c and the extreme ratio e, respectively. Each problem instance had $n_d = 5$ decision variables, $u = 1$ utility nodes and $r = 2$ variables were selected randomly as roots from the set of chance and decision variables. Each of the remaining chance and decision variables had $p = 2$ parents, while the cost function corresponding to the utility node had an arity of $a = 5$. The extreme probabilities were distributed according to the extreme ratio e which we varied between 0 and 0.95, respectively. The utility function contained only positive utility values which were generated uniformly randomly using the interval $[a = 0, b = 5]$ (see previous sections again for a refreshment). In all test cases we considered a sample set $\mathcal{E}_s$ of 100 decision policies drawn uniformly at random from the policy set of the respective order-of-magnitude influence diagram.

Figures 3(a)-(e) display the relative error RE^{OOM}_{avg} as a function of the problem size (which is given by the total number of variables), for five levels of the extreme ratio $e \in \{0, 0.25, 0.50, 0.75, 0.95\}$. Each data point in each of the plots represents the median value obtained for 10 random instances of the respective size. We conducted three sets of experiments for $\epsilon = 0.5$, $\epsilon = 0.05$ and $\epsilon = 0.005$, respectively. Informally, the smaller the ϵ value is, the more imprecise one is about the corresponding probability and utility values. Each experiment involved setting ϵ, evaluating the original influence diagram, translating the influence diagram into an order-of-magnitude influence diagram using the procedure from a previous section, and then evaluating the resulting order-of-magnitude influence diagram using the ELIM-OOM-ID algorithm.

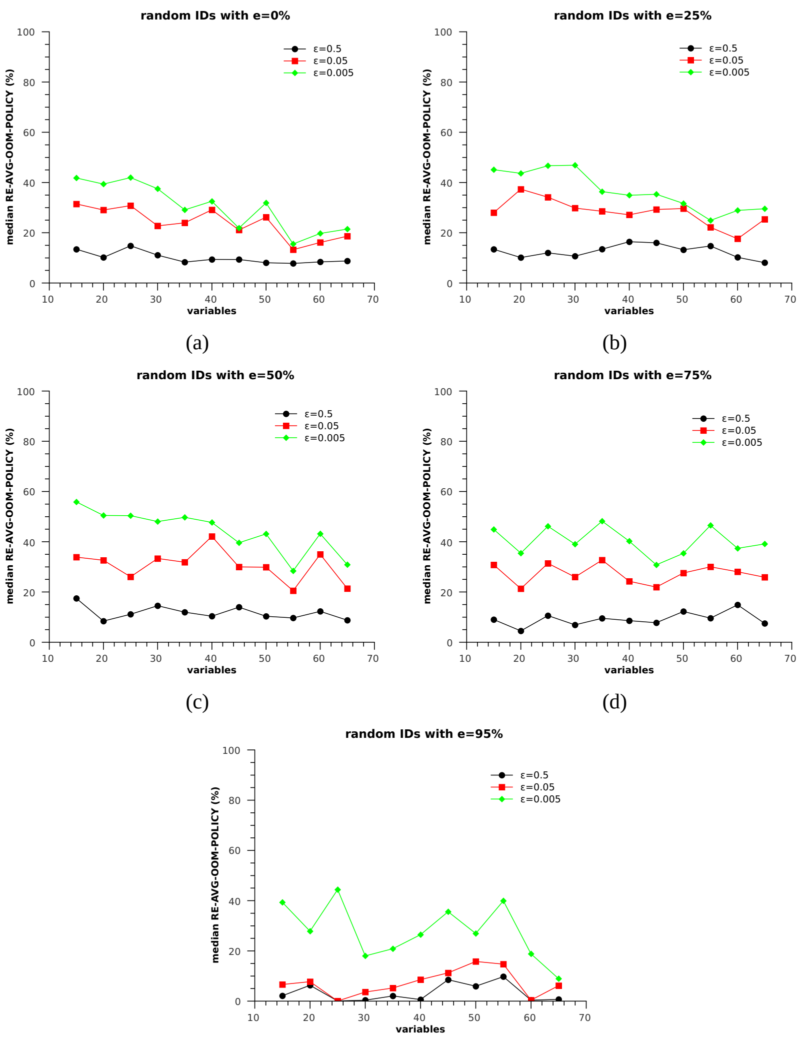

Figure 3: Results for influence diagrams with positive utility values. Shown is the relative error RE_{avg}^{OOM} of the OOM policies for $\epsilon \in \{0.5, 0.05, 0.005\}$ and extreme ratio $e \in \{0\%, 25\%, 50\%, 75\%, 95\%\}$. Average treewidth is between 8 and 23, respectively.

When $\epsilon = 0.5$, the relative error RE_{avg}^{OOM} is the smallest, ranging between 8% and 13% for $e = 0$ (see Figure 3(a)), between 8% and 16% for $e = 0.25$ (see Figure 3(b)), between 8% and 17% for $e = 0.50$ (see Figure 3(c)), between 7% and 14% for $e = 0.75$ (see Figure 3(d)) and between 0.06% and 9% for $e = 0.95$ (see Figure 3(e)), respectively.

When $\epsilon = 0.05$, the relative error RE_{avg}^{OOM} is larger, ranging between 13% and 31% for $e = 0$ (see Figure 3(a)), between 17% and 34% for $e = 0.25$ (see Figure 3(b)), between 20% and 42% for $e = 0.50$ (see Figure 3(c)), between 21% and 32% for $e = 0.75$ (see Figure 3(d)) and between 0.4% and 15% for $e = 0.95$ (see Figure 3(e)), respectively.

When $\epsilon = 0.005$, the relative error RE_{avg}^{OOM} is the largest, ranging between 15% and 41% for $e = 0$ (see Figure 3(a)), between 24% and 46% for $e = 0.25$ (see Figure 3(b)), between 28% and 55% for $e = 0.50$ (see Figure 3(c)), between 30% and 46% for $e = 0.75$ (see Figure 3(d)) and between 8% and 44% for $e = 0.95$ (see Figure 3(e)), respectively.

Note that the results obtained for the other two error measures, RE_{max}^{OOM} and RE_{med}^{OOM}, respectively, follow the same pattern (results to be included). Namely, the error was the smallest for $\epsilon = 0.5$ and it was the largest for $\epsilon = 0.005$.

Figures 4(a)-(e) display all four error measures, namely RE_{avg}^{OOM}, RE_{med}^{OOM}, RE_{min}^{OOM} and RE_{max}^{OOM}, as a function of the problem size for $e \in \{0, 0.25, 0.50, 0.75, 0.95\}$ and $\epsilon = 0.5$, respectively. We can see that RE_{max}^{OOM} is virtually zero especially for problems with a relatively large ratio of extreme probabilities (for $e > 0.5$ see Figures 4(c)-(e)). This means that the decision policy with maximum expected utility over the sample set $\mathcal{E}_s$ generated from the order-of-magnitude influence diagram for $\epsilon = 0.5$ was in most of the test cases identical to the optimal policy of the corresponding regular influence diagram. Furthermore, we can also see that on RE_{avg}^{OOM} stayed below 20% in all cases, which means that on average the order-of-magnitude decision policy was at most 20% off of the optimal policy of the corresponding regular influence diagram. Finally, we observed that for smaller values of ϵ, the results follow a similar pattern, however the errors were significantly larger than those obtained for the $\epsilon = 0.5$ case (results for $\epsilon = 0.05$ and $\epsilon = 0.005$ to be included as soon as they become available).

7.3.2 Influence diagrams with positive and negative utilities

For this experiment we generated random influence diagrams with the same parameters as before, except that in this case the utility values were generated uniformly randomly using the interval $[a = -5, b = 5]$, thus allowing for negative utility values as well (see again the model generator description for a refreshment).

Figures 5(a)-(e) display the relative error RE_{avg}^{OOM} as a function of the problem size (which is given by the total number of variables), for five levels of the extreme ratio $e \in \{0, 0.25, 0.50, 0.75, 0.95\}$. Each data point in each of the plots represents the median value obtained for 10 random instances of the respective size. As before, we have three sets of experiments for $\epsilon = 0.5$, $\epsilon = 0.05$ and $\epsilon = 0.005$, respectively.

As before, we can see that the relative error RE_{avg}^{OOM} is the smallest for $\epsilon = 0.5$ and it is the largest for $\epsilon = 0.005$, respectively. Note that in this case the differences between the different relative errors corresponding to different values of ϵ are

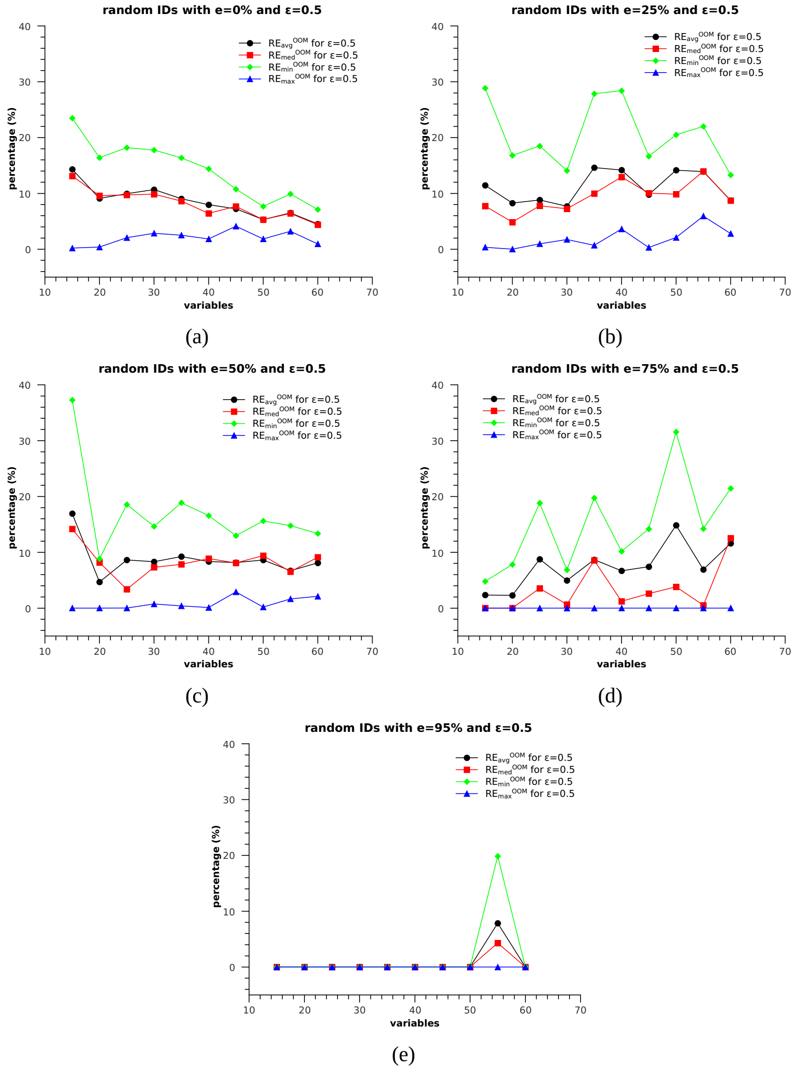

Figure 4: Results for influence diagrams with positive utility values. Shown are the relative errors RE^{OOM}_{avg}, RE^{OOM}_{med}, RE^{OOM}_{min} and RE^{OOM}_{max} of the OOM policies for $\epsilon = 0.5$ and extreme ratio $e \in \{0\%, 25\%, 50\%, 75\%, 95\%\}$. Average treewidth is between 8 and 23, respectively.

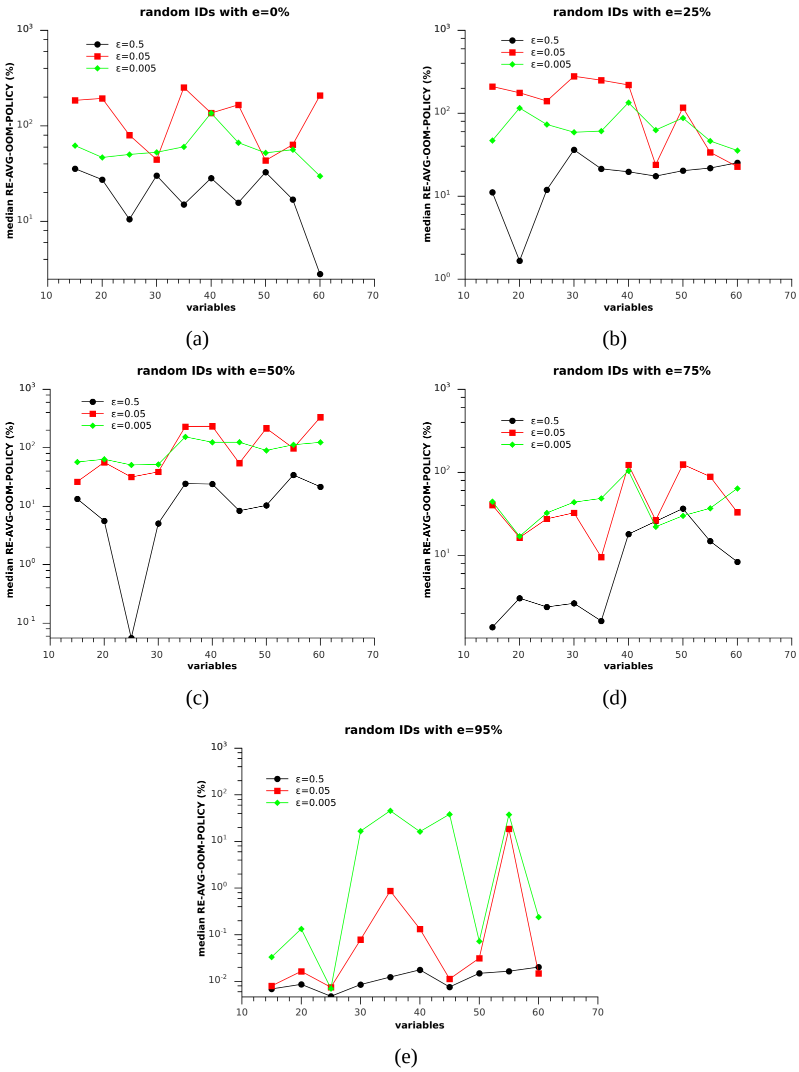

Figure 5: Results for influence diagrams with positive and negative utility values. Shown is the relative error of the OOM policies for $\epsilon \in \{0.5, 0.05, 0.005\}$ and extreme ratio $e \in \{0\%, 25\%, 50\%, 75\%, 95\%\}$. Average treewidth is between 8 and 23, respectively.

much more pronounced than those observed in the previous experiment, in many cases reaching two or more orders of magnitude (notice the logarithmic scale). For example, on problems with 35 variables and extreme ratio $e = 0.95$ (Figure 5(e)), the relative error RE_{avg}^{OOM} for $\epsilon = 0.5$ is about 2 and 4 orders of magnitude smaller than that for $\epsilon = 0.05$ and $\epsilon = 0.005$, respectively.

When $\epsilon = 0.5$, we can see that the relative error is virtually zero, especially for $e = 0.95$. In this case, RE_{avg}^{OOM} ranges between 0.006% and 0.02%, respectively (see Figure 5(e)). However, when the extreme ratio decreases, the relative error increases. Specifically, for $e = 0.75$ the error is between 1% and 36% (see Figure 5(d)), for $e = 0.50$ the error is between 0.05% and 24% (see Figure 5(c)), for $e = 0.25$ the error is between 1% and 25% (see Figure 5(b)) and for $e = 0$ the error is between 2% and 30% (see Figure 5(a)), respectively.

When $\epsilon = 0.05$ and $\epsilon = 0.005$, the relative error RE_{avg}^{OOM} is much higher spanning over several orders of magnitude as compared with the previous case ($\epsilon = 0.5$).

Figures 6(a)-(e) display all four error measures, namely RE_{avg}^{OOM}, RE_{med}^{OOM}, RE_{min}^{OOM} and RE_{max}^{OOM}, as a function of the problem size, for $e \in \{0, 0.25, 0.50, 0.75, 0.95\}$ and $\epsilon = 0.5$, respectively. As before, we can see that on problems with extreme ratio greater than 50%, RE_{max}^{OOM} is very close zero (see Figures 6(c)-(e)). When $e < 0.50$, RE_{max}^{OOM} is larger, reaching a as much as 14% on some problems with $e = 0.25$ (see Figure 6(b)). When looking at the average order-of-magnitude policies, we see that RE_{avg}^{OOM} is below 0.02% on problems with $e = 0.95$, and increases up to 30% on problems with $e \in \{0, 0.25, 0.50, 0.75\}$, respectively. This demonstrates again the robustness of the order-of-magnitude approximation for relatively large values of ϵ. When ϵ is small (typically less than 0.05), the relative error to the optimal policy of the corresponding regular influence diagram increases dramatically and therefore the quality of the order-of-magnitude policy degrades significantly.

7.3.3 Influence diagrams with negative utilities

Figures 7(a)-(e) display the relative error RE_{avg}^{OOM} as a function of the problem size (which is given by the total number of variables), for five levels of the extreme ratio $e \in \{0, 0.25, 0.50, 0.75, 0.95\}$. Each data point in each of the plots represents the median value obtained for 10 random instances of the respective size. As before, we have three sets of experiments for $\epsilon = 0.5$, $\epsilon = 0.05$ and $\epsilon = 0.005$, respectively.

Figures 8(a)-(e) display all four error meaures, namely RE_{avg}^{OOM}, RE_{med}^{OOM}, RE_{min}^{OOM} and RE_{max}^{OOM}, as a function of the problem size, for $e \in \{0, 0.25, 0.50, 0.75, 0.95\}$ and $\epsilon = 0.5$, respectively. As before, we can see that on problems with extreme ratio greater than 50%, RE_{max}^{OOM} is very close zero (see Figures 8(c)-(e)). On the other hand, RE_{avg}^{OOM} on average spans one or two orders of magnitude across all reported values of e and ϵ, respectively.

8 Conclusion and Future Work

The paper presents a new framework for qualitative sequential decision making under uncertainty based on an Order-of-Magnitude representation of probabilities and

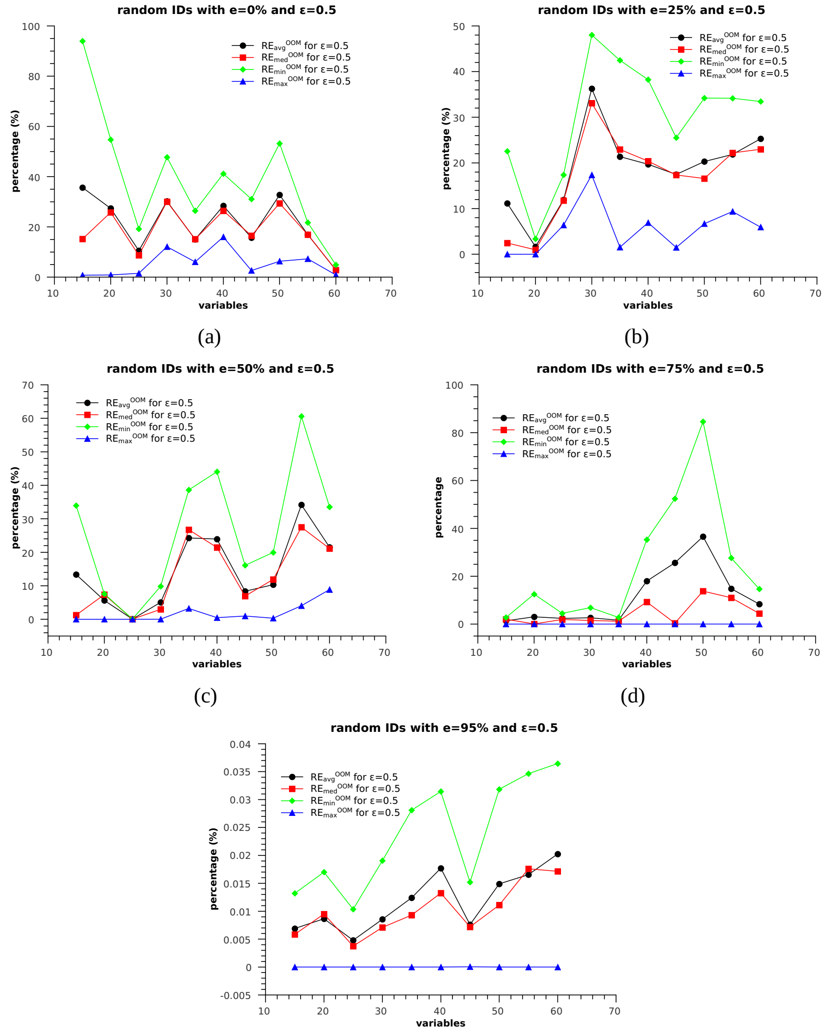

Figure 6: Results for influence diagrams with positive and negative utility values. Shown are the relative errors RE^{OOM}_{avg}, RE^{OOM}_{med}, RE^{OOM}_{min} and RE^{OOM}_{max} of the OOM policies for $\epsilon = 0.5$ and extreme ratio $e \in \{0\%, 25\%, 50\%, 75\%, 95\%\}$. Average treewidth is between 8 and 23, respectively.

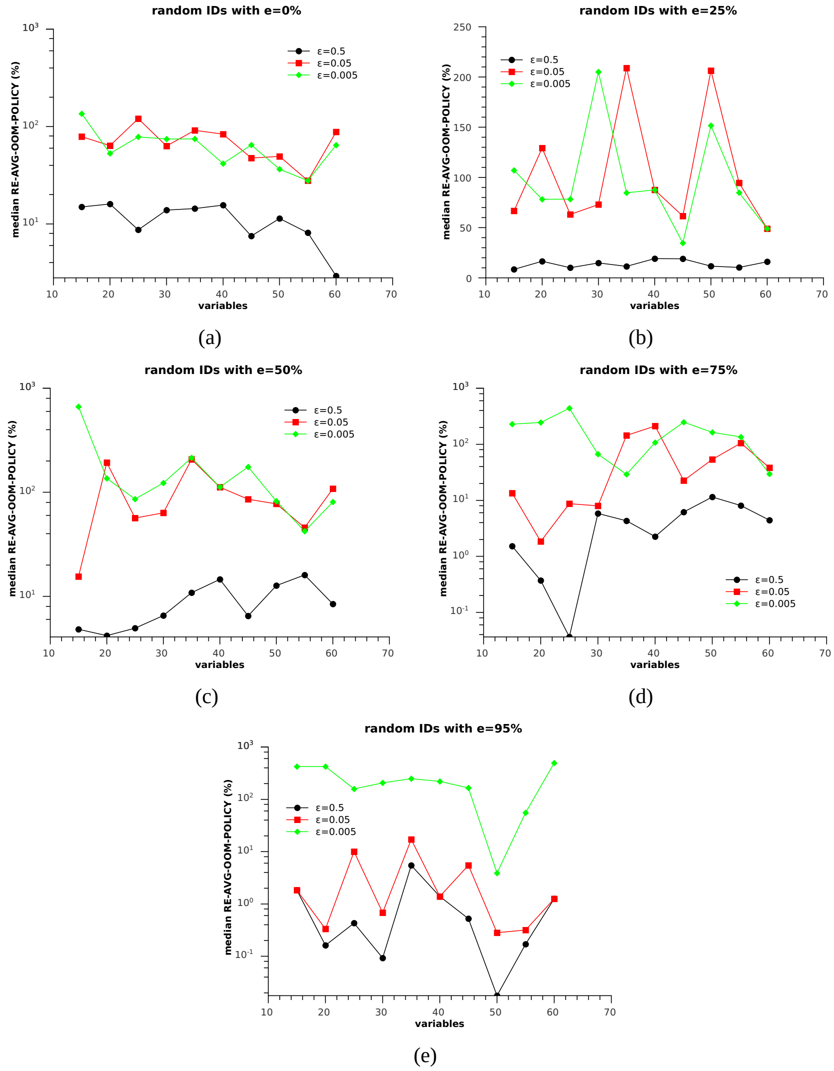

(a) (b) (c) (d) (e)

Figure 7: Results for influence diagrams with negative utility values. Shown is the relative error of the OOM policies for $\epsilon \in \{0.5, 0.05, 0.005\}$ and extreme ratio $e \in \{0\%, 25\%, 50\%, 75\%, 95\%\}$. Average treewidth is between 8 and 23, respectively.

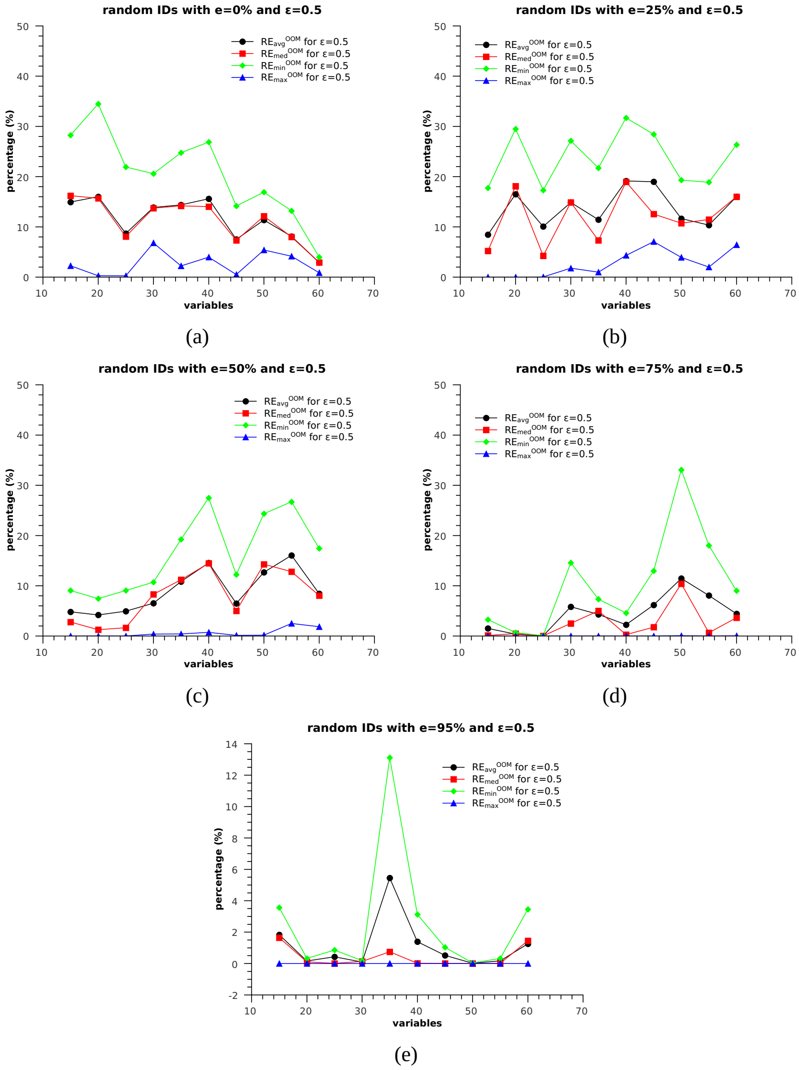

Figure 8: Results for influence diagrams with negative utility values. Shown are the relative errors RE^{OOM}_{avg}, RE^{OOM}_{med}, RE^{OOM}_{min} and RE^{OOM}_{max} of the OOM policies for $\epsilon = 0.5$ and extreme ratio $e \in \{0\%, 25\%, 50\%, 75\%, 95\%\}$. Average treewidth is between 8 and 23, respectively.

utilities. In particular, we introduce the Order-of-Magnitude Influence diagrams that extend the usual influence diagrams by replacing the point probability and utility values by order-of-magnitude probability and utility values, respectively. We also derive a sound variable elimination algorithm for computing an optimal policy that maximizes the order-of-magnitude expected utility. Numerical experiments on random influence diagrams show that in many cases the optimal policy of an order-of-magnitude influence diagram is almost identical to the optimal policy of a corresponding regular influence diagram.

Future work includes the computation of the optimal policy using depth-first or best-first heuristic search over a weighted AND/OR search graph associated with an order-of-magnitude influence diagram. In this direction we also plan to compile the policy set of an order of magnitude influence diagram into a multi-valued AND/OR decision diagram (AOMDD) to support sensitivity analysis tasks.

References

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

[2] P.P. Shenoy. Valuation-based systems for bayesian decision analysis. *Operations Research*, 40:463–484, 1992.

[3] F.V. Jensen F. Jensen and S.L. Dittmer. From influence diagrams to junction trees. In *Tenth Conference on Uncertainty in Artificial Intelligence*, pages 367–363, 1994.

[4] R. Dechter. A new perspective on algorithms for optimizing policies under uncertainty. In *International Conference on Artificial Intelligence Planning Systems (AIPS-2000)*, pages 72–81, 2000.

[5] R. A. Howard and J. E. Matheson. *Influence diagrams. The principles and applications of Decision analyis*. Strategic decisions Group, Menlo Park, CA, USA, 1984.

[6] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988.

[7] N. L. Zhang, R. Qi, and D. Poole. A computational theory of decision networks. *International Journal of Approximate Reasoning*, 11(2):83–158, 1994.

[8] M. Wellman. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, 44(1):257–303, 1990.

[9] W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. *Causation in Decision, Belief Change and Statistics*, 2(1):105–134, 1987.

[10] W. Spohn. A general non-probabilistic theory of inductive reasoning. In *Uncertainty in Artificial Intelligence*, pages 149–158, 1990.

[11] M. Goldszmidt. *Qualitative probabilities: a normative framework for commonsense reasoning*. PhD thesis, University of California, Los Angeles, 1992.

[12] L. Garcia and R. Sabbadin. Possibilistic influence diagrams. In *European Conference on Artificial Intelligence*, pages 372–376, 2006.

[13] N. Wilson. An order of magnitude calculus. In *Uncertainty in Artificial Intelligence*, pages 548–555, 1995.

[14] A. Darwiche and M. Goldszmidt. On the relation between kappa calculus and probabilistic reasoning. In *Uncertainty in Artificial Intelligence*, pages 145–153, 1994.