Practical Modeling of Bayesian Decision Problems - Exploiting Deterministic Relations

Anders L. Madsen
Hugin Expert A/S
Anders.L.Madsen@hugin.com

Kristian G. Olesen
Aalborg University
kgo@cs.auc.dk

Søren L. Dittmer
Systematic Software Engineering
dittmer@systematic.dk

Abstract

The widespread use of influence diagrams to represent and solve Bayesian decision problems is still limited by the inflexibility and rather restrictive semantics of influence diagrams. In this paper, we propose a number of extensions and adjustments to the definition of influence diagrams in order to make the practical use of influence diagrams more flexible and less restrictive. In particular, we describe how deterministic relations can be exploited to increase the flexibility and efficiency of representing and solving Bayesian decision problems. The issues addressed in this paper were motivated by the construction of a decision support system for mission management of unmanned underwater vehicles.

Introduction

Much effort has been devoted to the formalization of decision problems in order to provide automated support for their solution. Influence diagrams (Howard & Matheson 1981; Shachter 1986) have proven to be a vital candidate for the solution of sequential decision problems and the literature is rich on proposals that extend their expressive power (Shenoy 1992; Nielsen & Jensen 1999), improve their solution (Jensen, Jensen, & Dittmer 1994; Madsen 1999), and extend their functionality (Dittmer & Jensen 1997; Shachter 1999). It does seem, however, that influence diagrams have major obstacles in finding their way into widespread daily use. Experience shows that users have difficulties with the construction of influence diagrams and in particular with the semantics of information arcs. In the practical use of influence diagrams, it is a frequently occurring problem that the partial order imposed by the information arcs is difficult to predetermine. For diagnostic problems it is often difficult to foresee exactly which observations will be available at a given time. A similar problem arises in the case of unreliable sensors. By exploiting the deterministic nature of observations we point out that the strict order of calculations in the solution of influence diagrams can be relaxed. Deterministic relations, in general, can be utilized to improve the efficiency and flexibility of representing and solving decision problems.

The ideas presented in this paper rest on experience from ADVOCATE - an EU funded project. The aim of ADVOCATE is to design and develop new software architectures which add intelligence to existing software to cover different situations arising from any dysfunction of unmanned underwater vehicles (UUVs). ADVOCATE uses Bayesian networks and influence diagrams in diagnosis, risk assessment, and mission management decisions. Here we focus on mission management decisions related to missions involving the UUV system K-Fisch manufactured by STN-ATLAS Electronic. During the construction of an influence diagram representation of the mission management decision problem a number of interesting issues related to exploiting deterministic relations in the representation and solution of influence diagrams have emerged. The issues are described in terms of the influence diagram developed for the underwater vehicle mission management.

The main contribution of this paper is a deeper understanding of how to exploit deterministic relations when modeling Bayesian decision problems as influence diagrams in order to increase the efficiency and flexibility both representationally and computationally.

Preliminaries and Notation

An influence diagram is a triple $\mathrm{ID} = (\mathsf{G}, \mathcal{P}, \mathcal{U})$ where $\mathsf{G} = (\mathcal{V}, \mathcal{E})$ is a directed acyclic graph with chance, decision, and utility nodes, $\mathcal{P}$ is a set of conditional probability distributions, and $\mathcal{U}$ is a set of local utility functions. The nodes of G are connected by arcs. An arc into a chance node indicates a possible probabilistic dependence relation and an arc from a node X into a decision node D indicates that the state of X is known when decision D is to be made. Arcs into decision nodes are referred to as information arcs. Arcs into a utility node U specify the domain of the corresponding local utility function. To solve a decision problem is to determine an optimal strategy $\hat{\Delta} = \{\delta_i\}_{i \in [1;n]}$ consisting of a decision rule δ_i for each decision D_i and to compute the maximum expected utility $\hat{U}(\hat{\Delta})$ of $\hat{\Delta}$.

Let $\mathcal{R} = \{R_1, \ldots, R_m\}$ and $\mathcal{D} = \{D_1, \ldots, D_n\}$ denote the chance and decision variables, respectively. Furthermore, let $\mathcal{P} = \{P(R_1 \mid pa(R_1)), \ldots, P(R_m \mid pa(R_m))\}$ and $\mathcal{U} = \{u_1, \ldots, u_o\}$ be sets of probability distributions and local utility functions, respectively. From $\mathcal{P}$ and $\mathcal{U}$, $\hat{u}(\hat{\Delta})$ can be computed as:

$$\hat{u}(\hat{\Delta}) = \mathop{\bowtie}_{X \in \mathcal{V}} \left(\left(\prod_{i=1}^{m} P(R_i \mid pa(R_i)) \right) \sum_{j=1}^{o} u_j \right), \quad (1)$$

where $\bowtie$ is the generalized marginalization operator introduced by (Jensen, Jensen, & Dittmer 1994).

The set of chance variables $\mathcal{R}$ is partitioned into disjoint information sets $\mathcal{I}_0, \ldots, \mathcal{I}_n$ relative to the decision variables. The partition induces a partial ordering $\prec$ on the variables of the decision problem. The set of variables $\mathcal{I}_i$ observed between decisions D_i and D_{i+1} precedes D_{i+1} and succeeds D_i in the ordering: $\mathcal{I}_0 \prec D_1 \prec \mathcal{I}_1 \cdots \prec D_n \prec \mathcal{I}_n$. We will for ease of exposition assume a total order on the decision variables. Equation 1 can be computed using local computation by iteratively eliminating variables one at a time. Let Y be the next variable to eliminate and let $\Phi_Y = \{\phi \in \Phi \mid Y \in dom(\phi)\}$ and $\Psi_Y = \{\psi \in \Psi \mid Y \in dom(\psi)\}$ where Φ and Ψ are the current sets of probability and utility potentials, respectively. Calculate $\phi_Y^* = \bowtie_Y \prod_{\phi \in \Phi_Y} \phi$ and $\psi_Y^* = \bowtie_Y \prod_{\phi \in \Phi_Y} \phi \sum_{\psi \in \Psi_Y} \psi$ and update the sets Φ and Ψ as $\Phi^* = \Phi \cup \{\phi_Y^*\} \backslash \Phi_Y$ and $\Psi^* = \Psi \cup \left\{ \frac{\psi_Y^*}{\phi_Y^*} \right\} \backslash \Psi_Y$. An elimination order σ is a legal elimination order, if all variables in information set $\mathcal{I}_i$ are eliminated before D_i for all i. Thus, $\mathcal{I}_0$ is the set of variables initially observed and $\mathcal{I}_n$ is the set of variables never be observed or observed after the last decision has been made. A legal elimination order can be organized as a strong junction tree (Jensen, Jensen, & Dittmer 1994).

An observation on a variable R implies that the exact state r of R is known and is entered as an instantiation potential. A instantiation potential $f(R)$ is a deterministic function with domain R taking on the value 1 for the state r and 0 otherwise.

The ADVOCATE Project

ADVOCATE describes a distributed system for diagnosis and control of UUVs. The ADVOCATE system consists of a number of modules which communicate via a logical bus based on the CORBA standard, see figure 1. In this paper we focus on the mission management problem which involves monitoring the energy consumption in an autonomous underwater vehicle (AUV) and consequently the decisions on the further continuation of the mission.

The assumption is that at some point during operation of the AUV an unexpected reduction in the remaining energy will occur. At this point a decision on the further continuation of the mission has to be made by the mission management. An influence diagram model has been developed to support this decision. In order to support mission management decisions on the continuation of a mission is it necessary to represent the planned mission in the influence diagram. A decision on the continuation of the mission is based on the remaining energy and the estimated energy consumption of the underwater vehicle. The decision is either to proceed with the mission as planned, to abort the mission, or to decrease or increase the speed of the vehicle. A rough sketch of the model is shown in figure 2.

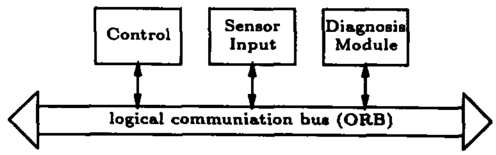

Figure 1: The CORBA architecture.

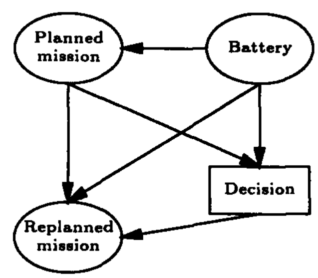

Figure 2: A rough sketch of the ADVOCATE model.

The influence diagram model was extracted from domain experts describing the physical relations between the components of the AUV, in particular the power supply. The uncertain relations stem from the uncertain measurements of the remaining energy and the exact discharge characteristics of the power supply as well as the expected energy consumption.

Uncertain Information Arcs

The information arcs of an influence diagram indicate the order in which information becomes available to the decision maker. During the construction of an influence diagram representation of a Bayesian decision problem, the model builder asks the domain expert to specify the order in which information becomes available. This may, however, not always be possible. For instance, a domain expert may want to specify that either she has made an observation on a particular variable R before the first decision, R will never observed, or R will be observed after the last decision has been made. In this case the set of variables initially observed ($\mathcal{I}_0$) is not known until the decision maker has specified which variables are observed. If the solution of an influence diagram is postponed until $\mathcal{I}_0$ and the states of the variables of $\mathcal{I}_0$ are known, then the computational efficiency can be improved w.r.t. both space and time. Usually, the variables of $\mathcal{I}_0$ are the last variables to be eliminated when solving an influence diagram ID, but since the solution of ID is postponed until $\mathcal{I}_0$ and the states of the variables of $\mathcal{I}_0$ are known, the variables of $\mathcal{I}_0$ can in fact be eliminated at any point.

The computational saving is caused by a reduction in the size of the effective junction trees. An effective junction tree is the junction tree used to perform the computations and not the junction tree constructed from the influence diagram representation of the problem. We propose to extend the definition of influence diagrams to include a notation - a dashed arc - specifying that a variable R either is observed prior to the first decision, after the last decision, or not at all. These arcs will be referred to as *uncertain information arcs*. The proposed extension can simplify the modeling task considerably as it facilitates the representation of different decision scenarios within the same influence diagram.

Example 1
Consider a decision problem involving an uncertain entity A which the decision maker may have two observations F_1 and F_2 on. This produces 4 different information scenarios, see figure 3. If F_1 and F_2 are known prior to the first decision D, then:

$$\hat{U}(\hat{\Delta}) = \sum_{F_1,F_2} \max_D \sum_A P(A)P(F_1|A)P(F_2|A)U(A,D).$$

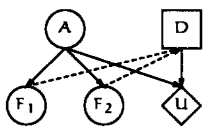

Figure 3: Observations on F_1 and F_2 are possible.

The deterministic relations induced by the instantiation potentials on F_1 and F_2 can be exploited to reduce the computational complexity of computing $\hat{U}(\hat{\Delta})$:

$$\hat{U}(\hat{\Delta}) = \max_D \sum_A P(A)U(A,D) \sum_{F_1} P(F_1|A)f(F_1) \sum_{F_2} P(F_2|A)f(F_2).$$

The difference in computational complexity of the above equations is even further emphasized, if the instantiations on F_1 and F_2 are entered by domain reduction rather than by marginalization. □

Eliminating the variables of $\mathcal{I}_0$ before any decision variable is eliminated will reduce the size of the junction tree. Furthermore, it is possible to exploit the same junction tree structure for a set of different decision scenarios. Due to the deterministic nature of instantiations, the variables of $\mathcal{I}_0$ can in fact be eliminated at any point during the solution of the influence diagram.

Proposition 1. *A deterministic potential* $f(X)$ *on a variable* $X \in \mathcal{I}_0$ *implies that* X *can be eliminated at any point.*

The fact that observed variables of $\mathcal{I}_0$ in decision problems with a single decision can be eliminated at any point has been used extensively when representing and solving influence diagrams. Corollary 2 shows that this property holds in general.

Corollary 2. *Observed variables of* $\mathcal{I}_0$ *can be eliminated at any point.*

Eliminating an instantiated chance variable $R \in \mathcal{I}_0$ before decision variable D_1 may produce decision rules with too large domains. That is, the domain of a decision rule δ_i ($\mathrm{dom}(\delta_i)$) for decision variable D_i may contain variables which are irrelevant for D_i. However, it is possible through a structural analysis to determine whether all variables of $\mathrm{dom}(\delta_i)$ are relevant for D_i, see e.g. (Nielsen & Jensen 1999). This can also be determined through a numerical analysis of δ_i.

Mission and Background Information

For the task of mission replanning it is necessary to represent information about the planned mission in the model. A planned mission is represented by the distance to cover (D_P) and the speed to cover the distance with (S_P). From D_P and S_P the time planned for the mission T_P is computed. At the point of decision, the progress of the mission (M) is known. This is used to compute the remaining distance (D_R) and the time used so far (T_U), see figure 4(a).

At the point of decision (D) information about the planned mission such as the planned speed (S_P), the progress of the mission (M), the remaining energy (R_M), and the planned distance to cover (D_P) may be known, see figure 4(b).

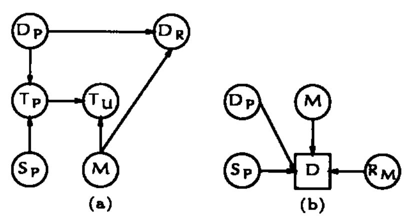

Figure 4: Planned mission(a) and the information available to the decision maker(b).

The time to complete the mission T_P is a deterministic function of D_P and S_P. This implies that $\hat{U}(\hat{\Delta})$ is independent of whether or not T_P is specified to be known by the decision maker before the decision.

In ADVOCATE readings on sensors and data in general might be lost. Thus, it is impossible for the domain experts to specify the set of background information variables. The solution of the decision problem is postponed until the set of observed variables is known. Thus, a single junction tree can be used to solve the decision problem for different decision scenarios.

Reduction of Possibilities

In a symmetric Bayesian decision problem the set of the decision options $\mathrm{dom}(D_i)$ available to the decision maker at a decision D_i are assumed to be independent of the past. This is, however, almost never the case in practice. It is often desirable to reduce the possible decision options at a decision in the future based on decisions and observations made in the past. The set of decision options available to the decision maker can be constrained by introducing a constraint variable C with two states y and n. The state of C is a deterministic function of its parent variables which returns y for all legal parent configurations and n otherwise. The legal configurations are enforced by instantiating C to y before the decision problem is solved.

Example 2
Assume that we are modeling a single player card game where the decision maker holds three cards (an ace (A), a king (K), and a queen (Q)). Assume the decision maker has to make two decisions. The first card to play (D_1) and the second card to play (D_2). Once a card has been played it cannot be played at a later point. Thus, the decision options at the second decision has to be constrained such that the same card is not played twice. This can be achieved by introducing a child constraint variable (C) of D_1 and D_2, see figure 5.

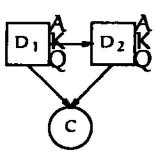

Figure 5: A single card can only be played once.

The influence diagram ID of figure 5 specifies that the state of C is not known until after both decisions have been made. This is, however, not the case. In fact, the state of C is known prior to the solution of ID. According to the traditional definition of influence diagrams the decision maker observes the variables of $\mathcal{I}_0$ and nothing else before the first decision is made. However, $C \notin \mathcal{I}_0$ and C is known before the first decision. Thus, according to the traditional definition of influence diagrams, ID is not a valid influence diagram due to the missing arcs from C to D_1 and D_2. However, due to the deterministic nature of constraint variables these arcs are not required. □

If some of the utilities are negative, then constraint variables can introduce serious problems as expected utilities might be zero ($EU = 0$) due to the utility being zero ($U = 0$) or due to zero probabilities ($p = 0$) representing illegal state space configurations. Thus, the use of constraint variables can imply that a decision option d for a decision variable D_i can have $EU(D_i = d) = 0$ while $EU(D_i = d') < 0$ for all $d' \neq d$. This, implies that $\max_{D_i} \sum_{\psi \in \Psi_{D_i}} \psi = d$, but selecting $D_i = d$ will produce an illegal state space configuration. Furthermore, the introduction of constraint variables makes the semantics of the model less clear. Notice that a linear transformation of the (total) utility function solves the problem of negative utilities.

We propose to model the asymmetry using deterministic functions specifying the legal configurations of decision variables based on the fact that constraint variables correspond to introducing a deterministic function $f(\mathcal{D}', \mathcal{R}')$ where $\mathcal{D}' \subseteq \mathcal{D}$ and $\mathcal{R}' \subseteq \mathcal{R}$. We are not introducing a new formalism to represent and solve asymmetric decision problems in general. Rather, we are proposing methods to handle restricted kinds of asymmetry within the framework of symmetric Bayesian decision problems and to solve these problems using existing efficient methods for solving symmetric decision problems. None of the existing methods for representing and solving asymmetric Bayesian decision problems (e.g. (Shenoy 2000) or (Nielsen & Jensen 2000)) are complete in the sense that all asymmetries can be represented without the introduction of special states and solved efficiently without some unnecessary computations. Therefore, it is still necessary to develop or extend the existing framework of influence diagrams for representing symmetric Bayesian decision problems to cope with asymmetries more efficiently.

Mission Termination

We want to model the fact that the mission should be aborted (A), if the mission cannot be completed given any of the possible speeds (D). That is, the mission should be aborted if the remaining energy (R_M) is less than the required energy (R_Q) to complete the mission given any speed, see figure 6.

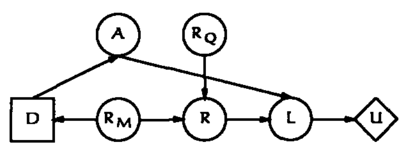

Figure 6: The mission is aborted if there is not enough energy to cover the distance given any speed.

The variable R specifies whether or not $R_M < R_Q$. The variable L specifies whether or not the vehicle is lost during the mission. To ensure the mission is aborted when $R_M < R_Q$, the vehicle is assumed to be lost if the mission is not aborted. This situation could be modeled more elegantly, if the possible decision options could be restricted using a deterministic function over D, R_Q, and R_M.

Deterministic Relations

In general, deterministic relations in influence diagrams are not only introduced by instantiations of variables and constraint variables. Deterministic relations can be specified in the conditional probability distributions of the model. Deterministic relations can be exploited to relax the constraints imposed by information arcs on the possible orders of elimination. Additional degrees of freedom in the selection of elimination orders can be exploited to increase the computational efficiency of solving Bayesian decision problems. A deterministic child variable X can be eliminated at any point in time up to the point where all its parent variables are known.

Example 3
Consider the influence diagram ID depicted in figure 7(a) and assume the relationship between S and T can be described as a deterministic function.

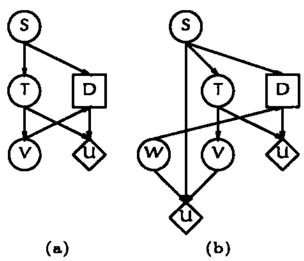

Figure 7: Two examples where the relationship between S and T is assumed to be deterministic.

Due to the deterministic relation between S and T, we can compute maximum expected utility $\hat{U}(\hat{\Delta})$ of the optimal strategy $\hat{\Delta}$ as:

$$\begin{aligned}\hat{U}(\hat{\Delta}) &= \sum_V \sum_S P(S) \max_D \sum_T P(T|S)P(V|T)U(T,D)\\ &= \sum_V \sum_S P(S) \sum_T \max_D P(T|S)P(V|T)U(T,D).\end{aligned}$$

That is, the marginalization of D and T can be commuted. Next, assume S = s is known which implies T = t and produces:

$$\begin{aligned}\hat{U}(\hat{\Delta}) &= \sum_V P(V|t)P(s)P(t|s) \max_D U(t,D)\\ &= \sum_V P(V|t)P(s) \max_D U(t,D).\end{aligned}$$

Thus, the deterministic relationship between S and T can be exploited both if the computations are performed before S is known, but especially if the computations are performed after S is known. □

Deterministic relations cannot only be exploited to commute the marginalization of chance and decision variables, they can also be exploited to distribute the marginalization of decision variables over local utility functions.

Example 4
Consider the influence diagram ID depicted in figure 7(b) and assume once again that the relationship between S and T can be described as a deterministic function. Exploiting the deterministic relationship between S and T, the maximum expected utility $\hat{U}(\hat{\Delta})$ of the optimal strategy $\hat{\Delta}$ is computed as:

$$\begin{aligned}\hat{U}(\hat{\Delta}) &= \sum_W P(W) \sum_S P(S) \max_D \sum_T P(T|S)\\ &\quad \Big(U(D,T) + \sum_V P(V|T)U(S,V,W)\Big) \qquad (2)\\ &= \sum_W P(W) \sum_S P(S) \sum_T \Big((\max_D P(T|S)U(D,T))\\ &\quad + P(T|S) \sum_V P(V|T)U(S,V,W)\Big). \qquad (3)\end{aligned}$$

Equation 3 does not offer a decrease in the total number of operations performed in order to solve ID compared to equation 2. It can, however, be used to reduce the number of multiplications and additions by increasing the number of maximizations. □

Battery Model

The estimated energy consumption (R_Q) to cover the remaining distance (D_R) given different speeds (S_R) is a function of the current (C) drawn from the battery and the remaining time (T_R), see figure 8(a). The speed S_R should be selected such that $R_M > R_Q$. In order to compute R_Q, it is necessary to represent the battery in the model. The power (P) drawn from the battery is computed as a function of S_R, the drag factor (D_F), the efficiency of the thrusters (E_T), the cross section (C_S), and the density of the water (D_W). The current C drawn from the battery is a function of P and the voltage (V), see figure 8(b).

The age and type of the battery and the water temperature are factors affecting the efficiency of the battery. As the age and type of the battery are not known for certain a probability distribution is specified.

Costs and Rewards

For the sake of completeness we briefly describe the cost and reward functions of the ADVOCATE model.

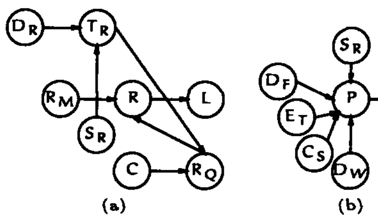

Figure 8: Energy consumption(a) and battery(b).

The utility function is decomposed into several local functions combining measures such as the cost of losing an AUV, cost of mission, cost of delays, reward from successful mission, e.t.c. The cost of a mission is in part determined by the total time used on the mission $T_T = T_R + T_U$ and whether or not the mission is completed with a delay (T_D), see figure 9. Besides the costs of a mission there are rewards of completing a mission.

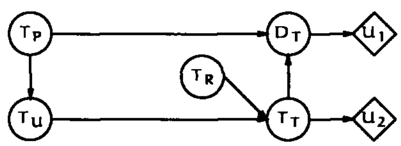

Figure 9: The cost of a mission is in part determined by the time used to complete the mission.

Discussion

Deterministic relations can be exploited to reduce the computational complexity of solving Bayesian decision problems and to ease the construction of influence diagram representations of Bayesian decision problems. The deterministic relations we have considered are induced by knowledge about the state of a variable, constraints on decision options, or deterministic conditional probability distributions. We have proposed extensions to the traditional definition of influence diagrams exploiting deterministic relations in order to increase the efficiency of representation and solution of Bayesian decision problem. Furthermore, deterministic relations between variables can, for instance, be used to simplify value of information analysis (Dittmer & Jensen 1997; Shachter 1999). Due to space constraints the applicability of deterministic relations cannot be described properly.

The extensions of representation and solution of influence diagrams proposed were motivated by the construction of a real-world decision support system for mission management of UUVs. At any given point in time of a mission the remaining energy of the battery can be measured and the progress of the mission is known. Thus, the influence diagram model can be used to continuously monitor the energy consumption. If the remaining energy of the battery is insufficient to complete the planned mission with the selected speed, the mission is replanned. The mission can be adjusted by either aborting the mission, or reducing or increasing the speed. The effectiveness of the mission management model constructed in the ADVOCATE project is currently being evaluated through the use of data gathered during previous missions with the K-Fisch UUV system. STN-ATLAS has found the model to be very useful for the mission management decision making. The usefulness of the model is emphasized by the intuitive structure and the close relationship between the model and the way decisions are made by mission management. The model decomposes into different subnetworks which naturally can be associated with entities of the modeled problem domain. This makes it particularly simple to apply the model in a different setting, i.e. to consider a different battery model.

Acknowledgement

This research was in part supported by ADVOCATE which is a project funded by the European Union under the Esprit Programme in the 4th Frame Programme (`http://advocate.e-motive.com`). We would like to thank Mr. P. Prendel and Mr. W. Hornfeld, STN ATLAS Elektronik, Bremen, Germany.

References

Dittmer, S. L., and Jensen, F. V. 1997. Myopic value of information for influence diagrams. In *Proc. 13th UAI*, 142–149.

Howard, R. A., and Matheson, J. E. 1981. Influence Diagrams. In *The Principles and Applications of Decision Analysis*, volume 2. chapter 37, 721–762.

Jensen, F.; Jensen, F. V.; and Dittmer, S. 1994. From Influence Diagrams to Junction Trees. In *Proc. 10th UAI*, 367–373.

Madsen, A. L. 1999. *All Good Things Come to Those Who Are Lazy - Efficient Inference in Bayesian Networks and Influence Diagrams Based on Lazy Evaluation.* Ph.D. Dissertation, Department of Computer Science, Aalborg University.

Nielsen, T. D., and Jensen, F. V. 1999. Welldefined decision scenarios. In *Proc. 15th UAI*, 502–511.

Nielsen, T. D., and Jensen, F. V. 2000. Representing and solving asymmetric Bayesian Decision Problems. In *Proc. 16th UAI*, –.

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

Shachter, R. 1999. Efficient Value of Information Computation. In *Proc. 15th UAI*, 594–601.

Shenoy, P. P. 1992. Valuation-Based Systems for Bayesian Decision Analysis. *OR* 40(3):463–484.

Shenoy, P. P. 2000. Valuation Network Representation and Solution of Asymmetric Decision Problems. *European Journal of Operational Research* 121(3):579–608.