Decision Making in Qualitative Influence Diagrams

Silja Renooij & *Linda C. van der Gaag*
Utrecht University, Department of Computer Science
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
e-mail: {silja,linda}@cs.ruu.nl

Abstract

The increasing number of knowledge-based systems that build on a *Bayesian belief network* or *influence diagram* acknowledge the usefulness of these frameworks for addressing complex real-life problems. The usually large number of probabilities and utilities required for their application, however, is often considered a major obstacle. The use of *qualitative abstractions* may to some extent remove this obstacle. *Qualitative belief networks* and associated algorithms have been developed before. In this paper, we address *qualitative influence diagrams* and outline an efficient algorithm for qualitative decision making.

1 Introduction

In the late 1980s, the framework of *Bayesian belief networks* was introduced for *reasoning with uncertainty* [Pearl 1988]. The framework provides a formalism for encoding a joint probability distribution on a set of statistical variables and offers algorithms for probabilistic inference. In practice, reasoning with uncertainty is often performed to support a decision maker in solving complex real-life problems. The belief-network framework in itself does not provide for *decision making under uncertainty*, as decision making involves not only knowledge of the uncertainties in a problem under study, but also knowledge of the decisions that are at a decision maker's disposal and of the desirability of their uncertain consequences. The framework of *influence diagrams* is tailored to decision making [Howard and Matheson 1981]. It provides a formalism for capturing the various types of knowledge involved in a decision problem and offers algorithms for computing preferred decisions. The framework is closely related to the belief-network framework; in fact, influence diagrams may be looked upon as enhanced belief networks.

The belief-network and influence-diagram frameworks have demonstrated their practicability in a wide range of problem domains. Experience shows, however, that the usually large number of probabilities and utilities required poses a major obstacle to their application [Druzdzel and Van der Gaag 1995]. Motivated by this experience, the framework of *qualitative belief networks* was introduced in the early 1990s by M.P. Wellman (1990).

A qualitative belief network abstracts from numerical probabilities by encoding *qualitative probabilistic relationships* among its variables. For reasoning with a qualitative belief network, an elegant algorithm is available from M.J. Druzdzel and M. Henrion (1993). As belief networks may be extended to influence diagrams, qualitative belief networks may be enhanced to *qualitative influence diagrams* [Wellman 1990]. A qualitative influence diagram abstracts from the numerical quantities involved in a decision problem under study by encoding *qualitative probabilistic* and *preferential relationships* among its variables.

Since their introduction, research has focused mainly on qualitative belief networks, with less attention for qualitative influence diagrams. As we consider decision making a valuable addition to reasoning with uncertainty, we re-introduce qualitative influence diagrams and outline a new algorithm for efficient *qualitative decision making*, that builds on Druzdzel and Henrion's algorithm for qualitative reasoning with uncertainty.

The paper is organised as follows. In Section 2 we review the belief-network and influence-diagram frameworks. In Section 3 qualitative belief networks are presented. In Section 4 we introduce qualitative influence diagrams; in addition, we outline our algorithm for qualitative decision making. In Section 5 we give some conclusions and directions for further research.

2 Belief networks and influence diagrams

The framework of *Bayesian belief networks* for reasoning with uncertainty is rooted in probability theory [Pearl 1988]. It offers a formalism for encoding a joint probability distribution on a set of statistical variables, in which information about independences is explicitly separated from numerical quantities.

A *belief network* consists of a qualitative part and an associated quantitative part. The qualitative part is a graphical representation of the independences holding among the variables in the encoded probability distribution. It takes the form of an *acyclic directed graph* G. Each node A in G represents a statistical variable that takes one of a finite set of values. We assume all variables to be binary, taking one of the values *true* and *false*; for abbreviation, we use a to denote $A = true$ and $\bar{a}$ to denote $A = false$. The arcs of G with each other model the independences among the represented variables. Informally, we take an arc $A \rightarrow B$ to represent an influential relationship between the variables A and B; the arc's direction marks B as the *effect* of the *cause* A. Absence of an arc between two nodes means that the corresponding variables do not influence each other directly and, hence, are (conditionally) independent.

Associated with the qualitative part of a belief network are numerical quantities from the encoded distribution. With each node A in G is associated a set of *conditional probabilities* $\Pr(A \mid \pi(A))$, describing the joint influence of values for the causes $\pi(A)$ of A on the probabilities of A's values. These sets of probabilities constitute the quantitative part of the network.

Example 1 Consider the belief network shown in Figure 1. The network represents a

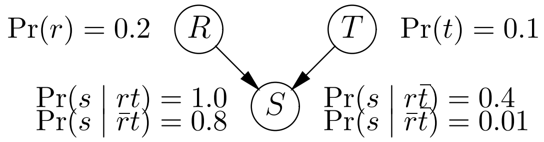

Figure 1: The *Sore Throat* belief network.

fragment of fictitious medical knowledge in pediatrics. Node S represents the presence or absence in a child of a severe sore throat, R represents the presence or absence of an upper respiratory tract infection, and T represents whether or not a child suffers from tonsillitis. Upper respiratory tract infection and tonsillitis are modelled as the possible causes of a sore throat. Note that the presence of any of these causes suffices to considerably increase the probability of a severe sore throat in a child. □

A belief network uniquely represents a probability distribution. It thus provides for computing any probability of interest. To this end efficient algorithms are available [Pearl 1988, Lauritzen and Spiegelhalter 1988].

A Bayesian belief network may be extended to an *influence diagram* to allow for *decision making under uncertainty* [Howard and Matheson 1981]. The formalism of influence diagrams provides for encoding not only a probability distribution on a set of variables, but also the decisions that a decision maker can take and the desirability of their uncertain consequences.

As a belief network, an influence diagram consists of a qualitative part and a quantitative part. The qualitative part again is an acyclic directed graph. Three different types of node are discerned. A node representing a statistical variable is termed a *chance node*; it is generally depicted as a circle. A *decision node* models a decision variable, representing the various decision alternatives that are at the decision maker's disposal; the node's value is under control of the decision maker. A decision node is depicted as a square. The third type of node is the *value node*. It represents the desirability of the consequences that may arise from the various decisions. There is only one value node and it does not have any outgoing arcs; it is depicted as a hexagon. The arcs between the chance nodes again encode the independences among the represented statistical variables. An arc from a decision node into a chance node expresses an influence on the represented statistical variable, exerted by the decision maker through a decision for the decision variable at hand. The incoming arcs of a decision node capture the information that is available before a decision is made. To conclude, an incoming arc of the value node expresses an influence on desirability.

The quantitative part of an influence diagram again associates with each chance node A in the diagram's digraph a set of conditional probabilities $\Pr(A \mid \pi(A))$. With the value node V is associated a set of *utilities* $u(\pi(V))$, specifying for each combination of values for V's parents $\pi(V)$ a number expressing the desirability of this value combination to the decision maker.

Example 2 Consider the influence diagram shown in Figure 2. The diagram embeds the

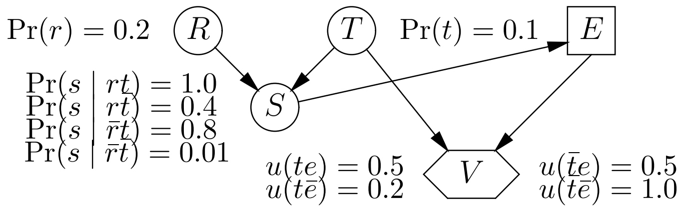

Figure 2: The *Sore Throat* influence diagram.

Sore Throat belief network from Figure 1. In addition, it includes the decision node E and the value node V. Node E models the decision alternatives that are at the decision maker's disposal; these are the decision to perform a tonsillectomy and the decision to refrain from performing one. A decision is made only if it is known with certainty whether or not a child is suffering from a severe sore throat. The preferred decision is to perform a tonsillectomy in the presence of tonsillitis and to refrain from performing a tonsillectomy in the absence of tonsillitis. □

An influence diagram uniquely represents a decision problem. A solution to the problem is a decision or, in case of multiple decision nodes, a sequence of decisions that *maximises* desirability of consequences. To compute a solution, for each sequence of decisions, the utilities of its uncertain consequences are weighted with the probabilities that these consequences will occur; the *expected utility* of the sequence x is thus computed from

$$\hat{u}(x) = \sum_i u(\pi_i(V)) \cdot \Pr(\pi_i(V) \mid x)$$

where $\pi_i(V)$ is a combination of values for the parents of the value node V and $u(\pi_i(V))$ is its utility; $\Pr(\pi_i(V) \mid x)$ is the probability of $\pi_i(V)$ given that the decisions x are taken. The preferred sequence of decisions is a sequence with highest expected utility. Efficient algorithms are available for decision making with influence diagrams [Shachter 1986].

3 Qualitative belief networks

Qualitative belief networks, introduced by M.P. Wellman as qualitative abstractions of belief networks, bear a strong resemblance to their quantitative counterparts [Wellman 1990]. A qualitative belief network comprises a graphical representation of the independences among a set of statistical variables, once more taking the form of an acyclic digraph. Instead of conditional probabilities, however, a qualitative belief network associates with its digraph *qualitative probabilistic relationships*.

A *qualitative influence* between two nodes expresses how the values of one node influence the probabilities of the values of the other node. For example, a *positive qualitative influence* of node A on its effect B, denoted $S^+(A, B)$, expresses that observing higher values for A makes higher values for B more likely, regardless of any other direct influence on B, that is,

$$\Pr(b \mid ax) \geq \Pr(b \mid \bar{a}x)$$

for any combination of values x for the set $\pi(B) \setminus \{A\}$ of causes of B other than A. A *negative qualitative influence*, denoted $S^-(A,B)$, and a *zero qualitative influence*, denoted $S^0(A,B)$, are defined analogously, replacing $\geq$ in the above formula by $\leq$ and $=$, respectively. If the influence of A on B is not monotonic, we say that it is *ambiguous*, denoted $S^?(A,B)$.

The set of influences of a qualitative belief network exhibits various convenient properties [Wellman 1990, Renooij 1996]. The property of *symmetry* guarantees that, if the network includes the influence $S^+(A,B)$, then it also includes $S^+(B,A)$. The property of *transitivity* asserts that qualitative influences along a chain, that specifies at most one incoming arc for each node, combine into a single influence with the $\otimes$-operator from Table 1. The property of *composition* asserts that multiple qualitative influences between two nodes along parallel chains combine into a single influence with the $\oplus$-operator. Note that combining qualitative influences may yield an ambiguous result. While for a qualitative influence along a single arc ambiguity indicates non-monotonicity, for a combined influence ambiguity may also indicate that its sign is unknown.

$\otimes$	$+$	$-$	0	$?$
$+$	$+$	$-$	0	$?$
$-$	$-$	$+$	0	$?$
0	0	0	0	0
$?$	$?$	$?$	0	$?$

$\oplus$	$+$	$-$	0	$?$
$+$	$+$	$?$	$+$	$?$
$-$	$?$	$-$	$-$	$?$
0	$+$	$-$	0	$?$
$?$	$?$	$?$	$?$	$?$

Table 1: The $\otimes$- and $\oplus$-operators.

In addition to influences, a qualitative belief network includes *synergies* modeling interactions among influences. An *additive synergy* between three nodes expresses how the values of two nodes jointly influence the probabilities of the values of the third node [Wellman 1990]. For example, a *positive additive synergy* of nodes A and B on their common effect C, denoted $Y^+(\{A,B\},C)$, expresses that the joint influence of A and B on C is greater than the sum of their separate influences, regardless of other influences on C, that is,

$$\Pr(c \mid abx) + \Pr(c \mid \bar{a}\bar{b}x) \geq \Pr(c \mid a\bar{b}x) + \Pr(c \mid \bar{a}bx)$$

for any combination of values x for the set $\pi(C) \setminus \{A,B\}$ of causes of C other than A and B. *Negative*, *zero*, and *ambiguous additive synergy* are defined analogously.

A *product synergy* expresses how the value of one node influences the probabilities of the values of another node in view of a given value for a third node [Henrion and Druzdzel 1991]; it describes an *intercausal* influence. For example, a *negative product synergy* of node A on node B (and vice versa) given the value c for their common effect C, denoted $X^-(\{A,B\},c)$, expresses that, given c, higher values for A render higher values for B less likely, that is,

$$\Pr(c \mid abx) \cdot \Pr(c \mid \bar{a}\bar{b}x) \leq \Pr(c \mid a\bar{b}x) \cdot \Pr(c \mid \bar{a}bx)$$

for any combination of values x for the set $\pi(C) \setminus \{A,B\}$. *Positive*, *zero*, and *ambiguous product synergy* again are defined analogously.

Example 3 We consider the qualitative abstraction of the *Sore Throat* belief network from Figure 1. From the conditional probabilities specified for node S, it is readily verified that both R and T exert a positive qualitative influence on S. As the joint influence of R and T on S is smaller than the sum of their separate influences, they exhibit a negative additive synergy on S. Furthermore, either value for node S induces an intercausal influence between R and T; this intercausal influence is described by a negative product synergy. The resulting qualitative belief network is shown in Figure 3. We would like to note that,

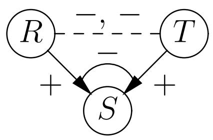

Figure 3: The qualitative *Sore Throat* belief network.

although in this example we have computed the qualitative probabilistic relationships from the probabilities of the original belief network, in real-life applications, these relationships are elicited directly from domain experts. □

For reasoning with a qualitative belief network, an elegant algorithm is available from M.J. Druzdzel and M. Henrion (1993). The basic idea of this algorithm is to trace the effect of observing a node's value on the other nodes in the network by message-passing between neighbouring nodes. For each node, a *sign* is determined, indicating the direction

procedure Propagate-Sign(*from*,*to*,*message*):

sign[*to*] $\leftarrow$ sign[*to*] $\oplus$ *message*;
for each (induced) neighbour V_i of *to*
do *linksign* $\leftarrow$ sign of (induced) influence
between *to* and V_i;
message $\leftarrow$ sign[*to*] $\otimes$ *linksign*;
if $V_i \neq$ *from* **and** $V_i \notin$ *Observed*
and sign[V_i] $\neq$ sign[V_i] $\oplus$ *message*
then Propagate-Sign(*to*,V_i,*message*)

of change in the node's probabilities occasioned by the new observation given all previously observed node values. Initially, all node signs equal '0'. For the newly observed node, an appropriate sign is entered, that is, either a '+' for the value *true* or a '−' for the value *false*. The node updates its sign and subsequently sends a message to each neighbour and every node on which it exerts an induced intercausal influence. The sign of this message is the $\otimes$-product of the node's (new) sign and the sign of the influence it traverses. This process is repeated throughout the network, building on the properties of symmetry, transitivity, and composition of influences. No node is visited more than twice.

4 Qualitative influence diagrams

Qualitative influence diagrams are qualitative abstractions of influence diagrams. A qualitative influence diagram, as its quantitative counterpart, comprises a representation of the variables involved in a decision problem along with their interrelationships, once more taking the form of an acyclic digraph. Instead of conditional probabilities, however, a qualitative influence diagram encodes qualitative influences and synergies on its chance variables. Instead of utilities, it specifies *qualitative preferential relationships*. These preferential relationships capture the preferences of the decision maker and, hence, pertain to the diagram's value node.

A *qualitative influence on utility* expresses how the values of a node influence expected utility. For example, a *positive qualitative influence on utility* of a parent A of the value node V, denoted $U^+(A)$, expresses that observing higher values for A increases expected utility, regardless of any other influence on utility, that is,

$$u(ax) \geq u(\bar{a}x)$$

for any combination of values x for the set $\pi(V) \setminus \{A\}$ of parents of V other than A. *Negative*, *zero*, and *ambiguous qualitative influences on utility* are defined analogously. As qualitative influences, influences on utility adhere to the properties of symmetry, transitivity, and composition; the symmetric counterpart of an influence on utility, however, is a qualitative influence and the transitive combination of a qualitative influence and an influence on utility is an influence on utility.

An *additive synergy on utility* expresses how the values of two nodes jointly influence expected utility. For example, a *positive additive synergy on utility* of two parents A and B of the value node V, denoted $Y_U^+(\{A, B\})$, expresses that the joint influence of the two nodes on expected utility is greater than the sum of their separate influences, that is,

$$u(abx) + u(\bar{a}\bar{b}x) \geq u(a\bar{b}x) + u(\bar{a}bx)$$

for any combination of values x for the set $\pi(V) \setminus \{A, B\}$. *Negative*, *zero*, and *ambiguous additive synergies on utility* are defined analogously. Note that as the value node of an influence diagram cannot be observed, product synergies on utility have no meaning.

Example 4 We consider the qualitative abstraction of the *Sore Throat* influence diagram from Figure 2. Since it embeds the qualitative belief network from Example 3, we focus on its preferential relationships. From the specified utilities, it is readily verified that node T exerts a negative qualitative influence on utility. The qualitative influence on utility of the decision node E is ambiguous as the desirability to the decision maker of a tonsillectomy depends on whether or not a child suffers from a tonsillitis. To conclude, T and E exhibit a positive additive synergy on utility. The resulting qualitative influence diagram is shown in Figure 4. □

For decision making with qualitative influence diagrams, M.P. Wellman has designed an algorithm based on the idea of recursively reducing a diagram [Wellman 1990]. Unfortunately, this algorithm tends to create more ambiguities than necessary and, hence, is able to

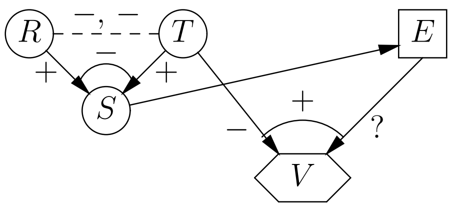

Figure 4: The qualitative *Sore Throat* diagram.

compute a preferred decision for fewer problems than possible [Druzdzel and Henrion 1993]. We propose a new algorithm for qualitative decision making, that builds on, and includes, Druzdzel and Henrion's algorithm for qualitative probabilistic inference. As the algorithm of Druzdzel and Henrion creates fewer ambiguities than Wellman's, our algorithm is able to solve more decision problems.

We recall that the algorithm of Druzdzel and Henrion traces the effect of observing a node's value on the other nodes in a qualitative belief network. As a qualitative influence diagram embeds a qualitative belief network for representing the relationships among its chance nodes, the algorithm can be applied straightforwardly to the diagram's probabilistic part. In addition, the algorithm can be used to trace the effect of the observation on the value node, yielding the sign of change in expected utility occasioned. The algorithm cannot be used, however, with regard to decision nodes, as it would ignore control of the decision maker.

To provide for decision nodes, we extend Druzdzel and Henrion's algorithm by including a second, similar process of message-passing between neighbouring nodes. This process is initiated by sending a '+' from the value node towards every decision node. Note that this message captures the decision maker's aim of maximising expected utility. The sign that

procedure Preferred-Decisions(*from*,*message*):

Propagate-Sign$_{\mathit{influence}}$(*from*,*from*,*message*);
Propagate-Sign$_{\mathit{utility}}$(V,V,'+');
for each decision node D
do if sign[*utility*,D] = '?' and $\alpha(D)$ causes the ambiguity
then sign[*utility*,D] $\leftarrow \oplus_i$ (sign[*influence*,A_i] $\otimes\ \delta_i$),
where $A_i \in \alpha(D)$ and δ_i is determined
from $Y_U^{\delta_i}(\{D, A_i\})$

thus reaches a decision node D reflects this node's influence on utility. So, if a '+' reaches D, the preferred decision is d; if a '$-$' reaches it, $\bar{d}$ is the preferred decision. If D receives a '0', then both decision alternatives are equally preferred. If D, however, receives an ambiguous sign, the preferred decision cannot be determined from the influence on utility of the node by itself. In fact, the ambiguity may indicate that the represented decision problem involves a true *trade-off*. By exploiting the signs of influence of the nodes that model the trade-off and their additive synergies on utility with node D, the ambiguity may be resolved; we illustrate the basic idea by means of our running example. Further details of our algorithm and a formal proof of its correctness will be provided in a forthcoming technical paper.

Example 5 Consider once more the qualitative *Sore Throat* influence diagram from Figure 4. Suppose that, after having observed a sore throat, we observe tonsillitis in a child. To reflect the new observation, a '+' is entered for node T. T updates its own sign to '+' and sends a '−' to nodes R and V; node R subsequently updates its sign of influence to '−'. Our algorithm now proceeds by sending a '+' from the value node V to the decision node E. Because of its ambiguous qualitative influence on utility, E receives a '?' and the preferred decision cannot yet be determined. From $U^?(E)$, we conclude, however, that either

$$u(te) > u(t\bar{e}) \text{ and } u(\bar{t}e) < u(\bar{t}\bar{e}), \text{ or}$$
$$u(te) < u(t\bar{e}) \text{ and } u(\bar{t}e) > u(\bar{t}\bar{e})$$

must hold. The first set of inequalities would correspond with a positive additive synergy on utility of nodes E and T, as it induces

$$u(te) + u(\bar{t}\bar{e}) \geq u(t\bar{e}) + u(\bar{t}e)$$

The second set of inequalities would correspond with a negative additive synergy on utility. Since the diagram specifies a positive additive synergy on utility of T and E, we know that the first set of inequalities holds. The preferred decision can now be determined: from the synergy, we have that in case of a positive sign of influence for T, the preferred decision is e, and in the case of a negative sign, the decision $\bar{e}$ is preferred. Since tonsillitis has actually been observed in the child under consideration, the algorithm yields the decision to perform a tonsillectomy as the preferred decision. □

5 Conclusions and further research

Qualitative abstractions of belief networks and influence diagrams have been introduced to remove the obstacle of acquiring a large number of probabilities and utilities. Research so far has focused mainly on qualitative belief networks. Since we consider decision making a valuable addition to reasoning with uncertainty, we have re-introduced the framework of *qualitative influence diagrams*. We have proposed a new algorithm for qualitative decision making under uncertainty, that builds on a similar algorithm for qualitative probabilistic reasoning. In developing our algorithm, we have assumed that a qualitative influence diagram under study includes binary variables only. Our algorithm is readily extended, however, to apply to more general diagrams.

One of the major drawbacks of qualitative abstractions is their coarse level of detail. Although for some problem domains this level will suffice, there are decision problems for which a finer level of detail is required. We would like to test our algorithm for qualitative decision making on various real-life applications to gain insight as to the level of detail generally required.

References

[Druzdzel and Van der Gaag 1995] Druzdzel, M.J. and Van der Gaag, L.C. 1995. Elicitation of probabilities for belief networks: combining qualitative and quantitative information. In Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence, 141 – 148.

[Druzdzel and Henrion 1993] Druzdzel, M.J. and Henrion, M. 1993. Efficient reasoning in qualitative probabilistic networks. In Proceedings of the Eleventh National Conference on Artificial Intelligence, 548 – 553.

[Henrion and Druzdzel 1991] Henrion, M. and Druzdzel, M.J. 1991. Qualitative propagation and scenario-based approaches to explanation in probabilistic reasoning. In Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence, 17 – 32.

[Howard and Matheson 1981] Howard, R.A. and Matheson, J. 1981. *Readings on the Principles and Applications of Decision Analysis*, Strategic Decisions Group, Menlo Park.

[Lauritzen and Spiegelhalter 1988] Lauritzen, S.L. and Spiegelhalter, D.J. 1988. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society, B*, 50: 157 – 224.

[Pearl 1988] Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Palo Alto.

[Renooij 1996] Renooij, S. 1996. Qualitative Probabilistic Networks: the Binary Case. Master's thesis, Dept. of Computer Science, Utrecht Univ.

[Shachter 1986] Shachter, R.D. 1986. Evaluating influence diagrams. *Operations Research*, 34: 871 – 882.

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