

Exploiting Non-monotonic Influences in Qualitative Belief Networks

Silja Renooij and Linda C. van der Gaag

Department of Computer Science, Utrecht University

P.O. Box 80.089, 3508 TB Utrecht, The Netherlands

{silja,linda}@cs.uu.nl

Abstract

In a qualitative belief network, dependences between variables are indicated by qualitative signs. These signs serve to model monotonic probabilistic relationships only: non-monotonic relationships between variables are modelled as lack of information. In this paper, we propose to include information about non-monotonic probabilistic influences between variables explicitly in a qualitative belief network. We show that this information can be exploited in probabilistic inference to forestall unnecessarily weak results.

Keywords: Qualitative belief networks, non-monotonic influences, probabilistic inference.

1 Introduction

In the late 1980s, the framework of *belief networks* was introduced for reasoning with uncertainty [4]. A belief network is a concise representation of a joint probability distribution on a set of statistical variables. It encodes the variables concerned along with their interdependences in a directed graph; the dependences between the variables are quantified by (conditional) probabilities. Associated with the belief-network formalism are algorithms for probabilistic inference. The increasing number of knowledge-based systems built on a belief network acknowledge the usefulness of the formalism and its associated algorithms for addressing complex real-life problems. Experience

shows, however, that the large number of probabilities required poses a major obstacle to their application [1]. Motivated by this experience, the framework of *qualitative belief networks* was introduced in the early 1990s [5].

A qualitative belief network is a qualitative abstraction of a quantified belief network. Like a belief network, it encodes the variables under consideration along with their interdependences in a directed graph. Rather than by probabilities, however, a qualitative belief network indicates the dependences between its variables by qualitative signs. These signs serve to capture qualitative probabilistic influences between variables. For probabilistic inference with a qualitative belief network, an elegant algorithm is available [2].

A qualitative belief network models, by means of signs, monotonic qualitative influences between its variables only. A qualitative influence between two variables is *monotonic* if observing higher values for one of these variables renders a shift in the probabilities of the values for the other variable in a direction that is not dependent upon any other influence. If the direction of shift does depend on influences from other variables, we say that the probabilistic influence between the two variables is *non-monotonic*.

In a qualitative belief network, a non-monotonic qualitative influence between two variables is indicated by a ‘?’ sign. The same sign is used to express an *unknown* qualitative influence, that is, a probabilistic influence for which the direction of shift is unknown. Non-monotonicity of a qualitative influence and lack of information therefore are expressed in the same way. Non-monotonicity

and lack of information, however, are different from a conceptual point of view. While an unknown qualitative influence does not provide any information, a non-monotonic influence conveys at least some information by the nature of its non-monotonicity. In this paper, we argue that it is worthwhile to explicitly distinguish between non-monotonic influences and unknown influences in a qualitative belief network. We show how useful information can be extracted from the non-monotonic influences of a network that can be exploited in probabilistic inference to forestall unnecessarily weak results.

The paper is organised as follows. In Section 2, we briefly review the belief-network framework; qualitative belief networks are introduced in Section 3. In Section 4, we investigate non-monotonic qualitative influences between variables and discuss how these influences can be exploited. The paper is rounded off with some conclusions and directions for further research in Section 5.

2 Belief networks

A *belief network* is a concise representation of a joint probability distribution on a set of statistical variables [4]. It consists of a qualitative part and an associated quantitative part. The qualitative part is a graphical representation of the interdependences between the variables in the encoded 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. In this paper, we assume all variables to be binary, taking one of the values *true* and *false*; for abbreviation, we use a to denote $A = \textit{true}$ and \bar{a} to denote $A = \textit{false}$. The arcs in the digraph G model possible dependences between the represented variables. Informally speaking, we take an arc $A \rightarrow B$ between the nodes A and B to represent an influential relationship between the associated 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 probability distribution. With each variable A in the digraph 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 variable A 's values. These sets of probabilities constitute the quantitative part of the network.

Example 1 We consider the belief network shown in Figure 1. The network represents

$$\begin{array}{ccc} \Pr(l) = 0.9 & \textcircled{L} & \textcircled{M} & \Pr(m) = 0.4 \\ & \searrow & \swarrow & \\ & \textcircled{C} & & \\ \Pr(c \mid lm) = 0.35 & & \Pr(c \mid l\bar{m}) = 0 & \\ \Pr(c \mid \bar{l}m) = 0.95 & & \Pr(c \mid \bar{l}\bar{m}) = 1.0 & \end{array}$$

Figure 1: The *Cervical Metastases* belief network.

a small, highly simplified fragment of medical knowledge in oncology, pertaining to lymphatic metastases of an oesophageal carcinoma. The variable L represents the location of an oesophageal carcinoma in a patient's oesophagus. The value *true* of L represents the information that the carcinoma resides in the lower two-third of the oesophagus; \bar{l} expresses that the carcinoma is located in the oesophagus' upper one-third. An oesophageal carcinoma upon growth typically gives rise to lymphatic metastases. The variable M represents the extent of these metastases. The value *false* of M indicates that just the local and regional lymph nodes are affected; m denotes that the distant lymph nodes are affected by cancer cells as well. Which lymph nodes are local or regional and which are distant depends on the location of the primary tumour in the oesophagus. The lymph nodes in the neck, or cervix, for example, are regional for a carcinoma in the upper one-third of the oesophagus and distant otherwise. The variable C represents the presence or absence in a patient of metastases in the cervical lymph nodes. \square

A belief network uniquely represents a joint probability distribution on its variables and thus provides for computing any probability of interest. Various different algorithms for probabilistic inference with a belief network are available.

3 Qualitative belief networks

Qualitative belief networks, as qualitative abstractions of belief networks, bear a strong resemblance to their quantitative counterparts [5]. A qualitative belief network equally comprises a graphical representation of the interdependences between a set of statistical variables, once again taking the form of an acyclic digraph. Instead of conditional probabilities, however, a qualitative belief network associates signs with its digraph. These signs serve to capture the probabilistic influences and synergies between the various variables.

A *qualitative probabilistic influence* between two variables expresses how the values of one variable influence the probabilities of the values of the other variable. For example, a *positive qualitative influence* of a variable A on its effect B , denoted $S^+(A, B)$, expresses that observing the value *true* for A makes the value *true* for B more likely, regardless of any other direct influences 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 non-monotonic, that is, the sign of the influence depends upon the values of other causes of B , or unknown, we say that the influence is *ambiguous*, denoted $S^?(A, B)$. With each arc in a qualitative network's digraph an influence is associated.

The set of influences of a qualitative belief network exhibits various convenient properties [5]. The property of *symmetry* guarantees that, if the network includes the qualitative influence $S^+(A, B)$, then it also includes $S^+(B, A)$. The property of *transitivity* asserts that the qualitative influences along a trail between two variables, specifying at most one incoming arc for each variable, combine into a single compound influence between these variables with the \otimes -operator from Table 1. The property of *composition* further asserts that multiple qualitative influences between

two variables along parallel trails combine into a compound influence between these variables with the \oplus -operator.

Table 1: The operators for combining signs.

\otimes	+	-	0	?	\oplus	+	-	0	?
+	+	-	0	?	+	+	?	+	?
-	-	+	0	?	-	?	-	-	?
0	0	0	0	0	0	+	-	0	?
?	?	?	0	?	?	?	?	?	?

In addition to influences, a qualitative belief network includes *synergies* modeling interactions between influences. An *additive synergy* between three variables expresses how the values of two variables jointly influence the probabilities of the values of the third variable. For example, a *positive additive synergy* of the variables 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 any other influences on C , that is,

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

for any combination of values x for the set of causes of C other than A and B . *Negative*, *zero*, and *ambiguous additive synergy* are defined analogously. A qualitative network specifies an additive synergy for each pair of causes and their common effect in its digraph.

A *product synergy* between three variables expresses how the value of one variable influences the probabilities of the values of another variable in view of an observed value for the third variable [3]. For example, a *negative product synergy* of a variable A on a variable B given the value *true* for their common effect C , denoted $X^-(\{A, B\}, c)$, expresses that, given c , the value *true* for A renders the value *true* for B less likely, that is,

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

for any combination of values x for the set of causes of C other than A and B . *Positive*, *zero*, and *ambiguous product synergy* again are defined analogously. For each pair of causes and their common effect, a qualitative belief network specifies two product synergies, one for each value of

the effect. Upon observation of a specific value for a common effect of two causes, the associated product synergy induces an influence between the two causes; the sign of this influence equals the sign of the synergy. A qualitative influence that is thus induced by a product synergy is termed an *intercausal influence*.

Example 2 We consider the qualitative abstraction of the *Cervical Metastases* belief network from Figure 1. From the conditional probabilities specified for the variable C , it is readily verified that the variable L exerts a negative qualitative influence on C ; the influence of the variable M on C is ambiguous. The joint influence of L and M on C is larger than the sum of their separate influences; L and M therefore exhibit a positive additive synergy on C . Furthermore, either value for the variable C induces an intercausal influence between L and M . For the value *true* of C this intercausal influence is captured by a positive product synergy and for the value *false* the influence is captured by a negative synergy. The qualitative belief network that is thus abstracted from the *Cervical Metastases* belief network is shown in Figure 2; the signs of the qualitative influences




Figure 2: The qualitative *Cervical Metastases* belief network.

are shown along the network’s arcs, the sign of the additive synergy is indicated over the curve over the variable C , and the signs of the product synergies are shown over the dashed line between the variables L and M . □

We would like to note that, although in the example above we have computed the signs of the various qualitative probabilistic relationships from the probabilities of the original belief network, in real-life applications these signs are elicited directly from domain experts.

For probabilistic inference with a qualitative belief network, an elegant algorithm is available [2]. The basic idea of this algorithm is to trace the effect of observing a single variable’s value on the

probabilities of the other variables represented in the network by message-passing between neighbours. For each variable, a *sign* is determined, indicating the direction of the shift in the variable’s probabilities occasioned by the new observation in the presence of previous observations. Initially, the sign of every variable equals ‘0’. For the newly observed variable, an appropriate sign is entered, that is, either a ‘+’ for the value *true* or a ‘-’ for the value *false*. The variable updates its sign and subsequently sends a message to each neighbour in the digraph and every variable on which it exerts an induced intercausal influence. The sign of this message equals the sign-product of the variable’s (new) sign and the sign of the influence associated with the arc it traverses. A variable that receives a message in turn updates its sign with the sign-sum of its original sign and the sign of the message it receives. If its sign has changed, the variable sends an appropriate message to any of its neighbours. This process is repeated throughout the network, building upon the properties of symmetry, transitivity, and composition of influences. Since a variable can change sign at most twice, the process visits each variable at most twice and is therefore guaranteed to halt.

4 Non-monotonic influences

A qualitative belief network serves to capture in essence only monotonic qualitative influences between its variables. We recall from Section 3 that, for example, a positive qualitative influence of a variable A on its effect B expresses that observing the value *true* for A makes the value *true* for B more likely. The influence exerted by A on B results in a shift in the probabilities of B ’s values in a direction that is independent of any other influences exerted on B . Qualitative influences between variables, however, need not necessarily be monotonic in nature as was demonstrated in Examples 1 and 2. The influence exerted by a variable A on its effect B is non-monotonic, for example, if the resulting direction of shift in the probabilities of B ’s values depends upon the influence of some other cause C on B .

In a qualitative belief network, a non-monotonic influence is denoted by the sign ‘?’’. The same sign is used to indicate an unknown qualitative

influence. Non-monotonicity of an influence and lack of information are thus represented in the same way. Non-monotonicity and lack of information, however, are not the same from a conceptual point of view. While an unknown qualitative influence does not provide any information at all, a non-monotonic influence conveys at least some information by the nature of its non-monotonicity. Now, the sign ‘?’ in a qualitative belief network gives rise to unwished-for ambiguous results in probabilistic inference, as is seen from Table 1. It is therefore worthwhile to try and avoid ‘?’-signs whenever possible. For this purpose, we will distinguish between the non-monotonic and unknown qualitative influences of a network explicitly and extract as much information as possible from its non-monotonic influences. We will show that this information can be exploited in probabilistic inference to forestall unnecessarily weak ambiguous results.

A *non-monotonic qualitative influence* of a variable A on its effect B is a qualitative influence of A on B that is not positive, negative, zero, or unknown. We say that the non-monotonicity of the influence is *provoked* by another cause C of B , denoted $S^{\sim c}(A, B)$, if the sign of the influence depends unambiguously on the value of C . More specifically, the non-monotonic influence expresses that for all combinations of values y for the set of causes of B other than A and C , we have either

$$\Pr(b | acy) \geq \Pr(b | \bar{a}cy) \quad \text{and} \\ \Pr(b | a\bar{c}y) \leq \Pr(b | \bar{a}\bar{c}y),$$

or

$$\Pr(b | acy) \leq \Pr(b | \bar{a}cy) \quad \text{and} \\ \Pr(b | a\bar{c}y) \geq \Pr(b | \bar{a}\bar{c}y),$$

with strict inequalities for at least one combination of values y . From this definition, it is readily seen that once a value for the provoking variable C has been observed, the non-monotonic influence of A on B reduces to a monotonic influence. We say that the observation *resolves* the non-monotonicity of A 's influence on B . We would like to note that the concept of provoking variable can easily be extended to sets of variables; for ease of exposition, however, we restrict the discussion to non-monotonicities provoked by a single variable.

Although an observation for its provoking variable reduces a non-monotonic influence between

two variables to a monotonic influence, the sign of the resulting influence is yet unknown. This sign, however, can be readily determined from the additive synergy defined for the variables concerned. We consider, as an example, a non-monotonic qualitative influence $S^{\sim c}(A, B)$ of a variable A on its effect B in which the non-monotonicity is provoked by the variable C . We suppose that the variables A and C exhibit a positive additive synergy on B , that is, we have

$$\Pr(b | acy) + \Pr(b | \bar{a}\bar{c}y) \geq \Pr(b | a\bar{c}y) + \Pr(b | \bar{a}cy)$$

for any combination of values y for the set of causes of B other than A and C . From the non-monotonicity of the influence of A on B and the sign of the additive synergy of A and C on B , we conclude that

$$\Pr(b | acy) \geq \Pr(b | \bar{a}cy) \quad \text{and} \\ \Pr(b | a\bar{c}y) \leq \Pr(b | \bar{a}\bar{c}y)$$

for any combination y . Now, upon observation of the value *true* for the provoking variable C , we find that

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

for any combination of values x , including the observation c , for the set of causes of B other than A . We conclude that, after resolving the non-monotonicity involved, the variable A exerts a positive qualitative influence on B . Alternatively, upon observation of \bar{c} , the variable A exerts a negative influence on B . A negative additive synergy of A and C on B leads to an analogous result. We conclude that the sign of the resolved non-monotonic influence equals the sign-product of the sign of the additive synergy involved and the sign of the observation for the provoking variable.

Example 3 We consider once again the qualitative *Cervical Metastases* belief network from Figure 2. The ambiguous influence $S^? (M, C)$ of the extent of the lymphatic metastases of an oesophageal carcinoma, M , on the presence of metastases in the cervical lymph nodes, C , is a non-monotonic influence in which the non-monotonicity is provoked by the location of the carcinoma, L . From the probabilities specified for the variable C in Figure 1, it is readily seen that, given a carcinoma in the upper one-third of

a patient's oesophagus, that is, given \bar{l} , the variable M exerts a negative influence on C ; given l , M exerts a positive qualitative influence on the variable C . In the qualitative belief network, the sign of the influence of M on C after resolution by \bar{l} is computed to be the sign-product of the sign '+' of the additive synergy of M and L on C and the sign '-' of the observation for L , that is, the sign is computed to be $+\otimes-=-$. After resolution by l , the sign of the influence of M on C is computed to be $+\otimes+=+$. \square

We would like to note that in real-life applications of a qualitative belief network, non-monotonic qualitative influences and their provoking variables are elicited directly from domain experts.

For probabilistic inference with a qualitative belief network in which non-monotonic and unknown influences are explicitly distinguished, basically the same algorithm can be used as for probabilistic inference with a regular qualitative network. The only difference lies in the traversal of a non-monotonic qualitative influence. Before propagating the sign of a non-monotonic influence by sign multiplication, it is investigated whether or not the influence's non-monotonicity is resolved by the available observations. If the non-monotonicity is resolved, the sign of the resolved influence as described above is propagated; otherwise, the ambiguous sign '?' is propagated. We would like to point out that by thus exploiting information about non-monotonic influences in a qualitative belief network, at least some ambiguous results in probabilistic inference are forestalled.

5 Conclusions and further research

A qualitative belief network in essence serves to capture monotonic probabilistic influences between its variables only. We have argued that it is worthwhile to explicitly capture information about non-monotonic influences as well. We have shown that this information can be exploited in probabilistic inference to forestall unnecessarily weak ambiguous results. In this paper, we have focused attention on non-monotonic influences between binary variables. We would like to extend our ideas to induced intercausal influ-

ences that are non-monotonic in nature. Furthermore, we would like to generalise our results to non-binary variables. To conclude, we envision further investigation of the information that can be derived from sets of variables provoking non-monotonicities to forestall even more ambiguous results in probabilistic inference with a qualitative belief network.

References

- [1] M.J. Druzdzel and L.C. van der Gaag (1995). Elicitation of probabilities for belief networks: combining qualitative and quantitative information. In: *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pp. 141 – 148.
- [2] M.J. Druzdzel and M. Henrion (1993). Efficient reasoning in qualitative probabilistic networks. In: *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pp. 548 – 553.
- [3] M. Henrion and M.J. Druzdzel (1991). Qualitative propagation and scenario-based approaches to explanation in probabilistic reasoning. In: *Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence*, pp. 17 – 32.
- [4] J. Pearl (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Palo Alto.
- [5] M.P. Wellman (1990). Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, vol. 44, pp. 257 – 303.