

Introducing Situational Signs in Qualitative Probabilistic Networks

Janneke H. Bolt

Linda C. van der Gaag

Silja Renooij

institute of information and computing sciences, utrecht university

technical report UU-CS-2004-006

www.cs.uu.nl

Introducing Situational Signs in Qualitative Probabilistic Networks

Janneke H. Bolt, Linda C. van der Gaag, Silja Renooij
Institute of Information and Computing Sciences, Utrecht University
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
{janneke,linda,silja}@cs.uu.nl

Abstract

A qualitative probabilistic network is a graphical model of the probabilistic influences among a set of statistical variables in which each influence is associated with a qualitative sign. A non-monotonic influence between two variables is associated with the ambiguous sign '??', which indicates that the actual sign of the influence depends on the state of the network. The presence of such ambiguous signs is undesirable as it tends to lead to uninformative results upon inference. In this paper, we argue that, in each specific state of the network, the effect of a non-monotonic influence is unambiguous. To capture the current effect of the influence, we introduce the concept of *situational sign*. We show how situational signs can be used upon inference and how they are updated as the state of the network changes. By means of a real-life qualitative network in oncology, we show that the use of situational signs can effectively forestall uninformative results upon inference.

1 Introduction

The formalism of Bayesian networks [1], is generally considered an intuitively appealing and powerful formalism for capturing the knowledge of a complex problem domain along with its uncertainties. A Bayesian network consists of a directed acyclic graph in which each node represents a statistical variable and each arc expresses a probabilistic relationship between the connected variables. To capture the strengths of these relationships, each variable has associated a set of conditional probability distributions that describe the effect of all possible combinations of values for its predecessors in the digraph, on the probabilities of its values. For reasoning with a Bayesian network, powerful algorithms are available.

For probabilistic reasoning in a qualitative way, qualitative abstractions of Bayesian networks were introduced. These *qualitative probabilistic networks* (QPNs) [2] equally encode statistical variables and the probabilistic relationships between them in a directed acyclic graph. The relationships between the variables are not quantified by conditional probabilities, however, but are summarised by qualitative signs instead. For reasoning with a qualitative probabilistic network in a mathematically correct way, an efficient algorithm is available that is based on the idea of propagating and combining these signs [3].

Recently, a methodology for building Bayesian networks was introduced in which the construction and validation of a qualitative network is proposed as an intermediate step [4]. For Bayesian networks, the usually large number of probabilities required, tends to pose a major obstacle to their construction [5]. By first building a qualitative network, with the help of domain experts, the reasoning behaviour of the Bayesian network in the making can be studied and validated prior to probability assessment. The signs of the validated qualitative network, moreover, can be used as constraints on the probabilities to be obtained for the Bayesian network, thereby simplifying the quantification task.

To exploit the use of qualitative probabilistic networks in a methodology for building Bayesian networks, inference with a qualitative network should yield results that are as informative as pos-

sible. Qualitative networks, however, model the probabilistic relationships between their variables at a higher abstraction level than Bayesian networks. Reasoning with a qualitative network can, as a consequence, lead to results that are not informative with respect to the Bayesian network in the making. In the past decade, various researchers have addressed this tendency of qualitative networks to yield uninformative results upon inference, and have proposed extensions to the basic formalism. Renooij and Van der Gaag [6], for example, have enhanced qualitative networks by adding a notion of strength to provide for resolving trade-offs between conflicting influences.

Closely linked with the high abstraction level of representation in qualitative probabilistic networks, is the issue of non-monotonicity. An influence of a variable A on a variable B is called non-monotonic if it is positive in one state and negative in another state of the network under consideration. A non-monotonic influence cannot be assigned an unambiguous sign of general validity and is associated with the uninformative ambiguous sign '?'. Although a non-monotonic influence may have varying effects, in each particular state of the network its effect is unambiguous. In this paper, we extend the framework of qualitative probabilistic networks with signs that capture information about the current effect of non-monotonic influences. These signs are termed *situational* to express that they are valid only in particular situations, or states of the network. We show how situational signs can be used upon inference and how they may forestall uninformative results. Because situational signs are associated with a particular state of the network under consideration, they are dynamic in nature and may need updating as the state of the network changes. We present a method for this purpose and adapt the standard algorithm for qualitative probabilistic reasoning to provide for inference with situational qualitative networks.

To investigate the practicability of situational signs, we study the effect of their introduction into a real-life qualitative network in the field of oesophageal cancer. We study the difference in performance between the original network with ambiguous signs for its non-monotonic influences and the same network in which these ambiguous signs have been supplemented with situational signs. For our study, we use the medical records of 156 patients diagnosed with cancer of the oesophagus. We show that the introduction of situational signs serves to decrease the percentage of ambiguities that are propagated to one of the main diagnostic variables in the network from 45% to 12%. More in general, the results of our study demonstrate that a network supplemented with situational signs can yield more informative results upon inference than the original qualitative network.

The paper is organised as follows. Section 2 provides some preliminaries on qualitative probabilistic networks and introduces the basic algorithm for reasoning with a qualitative network. Section 3 introduces the concept of situational sign. The dynamics of situational signs are detailed in Section 4, which also gives an adapted algorithm for reasoning with a situational qualitative network. In Section 5 the effect of introducing situational signs upon the performance of the qualitative oesophageal cancer network is described. The paper ends with our concluding observations in Section 6.

2 Preliminaries

In Section 2.1 we describe the formalism of qualitative probabilistic networks and in Section 2.2 we review the basic algorithm for probabilistic reasoning with a qualitative network.

2.1 Qualitative Probabilistic Networks

In this paper, we are concerned with probability distributions over sets of statistical variables. The (sets of) variables will be denoted by upper-case letters. We assume all variables A to be binary, taking one of the values a_1 and a_2 . We further assume that the values of A are ordered, with $a_1 > a_2$. For abbreviation, we write a to denote $A = a_1$, and \bar{a} to denote $A = a_2$.

As qualitative probabilistic networks were introduced as qualitative abstractions of Bayesian networks, we begin by briefly reviewing the latter type of network. A *Bayesian network* is a graphical model of a joint probability distribution Pr over a set of statistical variables. This




Figure 1: An example Bayesian network, modelling the effects of fitness (F) and training (T) on a feeling of well-being (W).

model comprises a directed acyclic graph in which each node represents a statistical variable. As there is a one-to-one correspondence between nodes and variables, we will use the two terms interchangeably. The probabilistic relationships between the variables are captured by the digraph's set of arcs. The absence of an arc $A \rightarrow B$ between the nodes A and B indicates that there is no direct influence between the associated variables. If all trails between A and B are *blocked* by the available evidence, moreover, there is also no indirect influence between the variables. We say that a trail between A and B is blocked by the evidence if it includes either an observed variable with at least one outgoing arc, or an unobserved variable with two incoming arcs and no observed descendants. If there is no influence between A and B , then the associated variables are conditionally independent given the available evidence [1].

Associated with the digraph of a Bayesian network are numerical probabilities from the modelled distribution. With each variable A is associated a set of (conditional) probability distributions $\Pr(A | \pi(A))$; each of these distributions describes the joint effect of a specific combination of values for the parents $\pi(A)$ of A on the probabilities of A 's values.

Example 1 We consider the small Bayesian network from Figure 1. The network represents a fragment of fictitious knowledge about the effect of training and fitness on one's feeling of well-being. Variable F captures one's fitness and variable T models whether or not one has undergone a training session; variable W models whether or not one has a feeling of well-being. All variables are binary, with the values *true* $>$ *false*. \square

In its initial state where no observations have been entered, a Bayesian network captures a prior joint probability distribution over its variables. When observations are entered for one or more of the variables discerned, the network converts to another state and then represents the posterior distribution given the entered evidence.

Since *qualitative probabilistic networks* are qualitative abstractions of Bayesian networks, they bear a strong resemblance to their quantitative counterparts. A qualitative network also comprises a directed acyclic graph modelling variables and the probabilistic relationships between them. Instead of conditional probability distributions, however, a qualitative probabilistic network associates with its digraph *qualitative influences* and *qualitative synergies* to capture features of the represented distribution [2].

A *qualitative influence* between two variables expresses how observing a value for the one variable affects the probability distribution over the values of the other variable. For example, a *positive qualitative influence* of a variable A on a variable B along an arc $A \rightarrow B$ expresses that observing the higher value for A makes the higher value for B more likely, regardless of any other direct influences on B , that is,

$$\Pr(b | ax) - \Pr(b | \bar{a}x) \geq 0$$

for any combination of values x for the set $\pi(B) \setminus \{A\}$ of parents of B other than A . The influence is denoted $S^+(A, B)$; the '+' is termed the *sign* of the influence. A negative qualitative influence, denoted S^- , and a zero qualitative influence, denoted S^0 , are defined analogously, replacing \geq in the above formula by \leq and $=$, respectively. For a positive, negative or zero influence of A on B , the difference $\Pr(b | ax) - \Pr(b | \bar{a}x)$ has the same sign for all combinations of values x . These influences thus describe a *monotonic* effect of a change in A 's distribution on the probability distribution over B 's values. If the influence of A on B is positive for one combination x and

negative for another combination, however, we say that the influence is *non-monotonic*. Non-monotonic influences are associated with the sign '??', indicating that their effect is ambiguous. The same sign is used for influences whose sign is unknown.

The set of all influences of a qualitative probabilistic network exhibits various important properties [2]. The property of *symmetry* states that, if the network includes the influence $S^\delta(A, B)$, then it also includes $S^\delta(B, A)$, $\delta \in \{+, -, 0, ?\}$. The *transitivity* property asserts that the qualitative influences along a trail that specifies at most one incoming arc for each variable, combine into a net influence whose sign is defined by the \otimes -operator from Table 1. The property of *composition* asserts that multiple influences between two variables along parallel trails combine into a net influence whose sign is defined by the \oplus -operator. The three properties with each other provide for establishing the sign of the net influence between any two variables in a qualitative network.

Table 1: The \otimes - and \oplus -operators for combining signs.

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

In addition to influences, a qualitative probabilistic network includes synergies to capture the joint interactions among three or more variables. An *additive synergy* between three variables expresses how the values of two variables interact to yield a joint effect on the probability distribution for the third variable. For example, a *positive additive synergy* of the variables A and C on their common child B , denoted $Y^+(\{A, C\}, B)$, expresses that A and C serve to strengthen each other's influence on B regardless of any other direct influences on B , that is,

$$\Pr(b \mid acx) + \Pr(b \mid \bar{a}\bar{c}x) \geq \Pr(b \mid a\bar{c}x) + \Pr(b \mid \bar{a}cx)$$

for any combination of values x for the set $\pi(B) \setminus \{A, C\}$ of parents of B other than A and C . A negative additive synergy, denoted Y^- , and a zero additive synergy, denoted Y^0 , are defined analogously. A non-monotonic or unknown additive synergy of variables A and C on B is denoted $Y^?(\{A, C\}, B)$. A *product synergy* expresses how the value of one variable influences the probability distribution over the values of another variable in view of a previous observation for a third variable [7]. As in this paper we will only consider single observations or multiple simultaneous observations, we will not further elaborate on this type of synergy.

Example 2 We consider again the Bayesian network from Figure 1 and construct its qualitative abstraction. From the conditional probability distributions specified for the variable W , we have that $\Pr(w \mid ft) - \Pr(w \mid \bar{f}\bar{t}) \geq 0$ and $\Pr(w \mid f\bar{t}) - \Pr(w \mid \bar{f}t) \geq 0$. We conclude that $S^+(F, W)$: fitness favours a feeling of well-being regardless of training. We further have that $\Pr(w \mid ft) - \Pr(w \mid f\bar{t}) > 0$ and $\Pr(w \mid \bar{f}t) - \Pr(w \mid \bar{f}\bar{t}) < 0$. We thus find that $S^?(T, W)$: the effect of training on well-being depends on one's fitness. From $\Pr(w \mid ft) + \Pr(w \mid \bar{f}\bar{t}) \geq \Pr(w \mid f\bar{t}) + \Pr(w \mid \bar{f}t)$, to conclude, we find that $Y^+(\{F, T\}, W)$. The resulting qualitative network is shown in Figure 2; the signs of the qualitative influences are shown along the arcs, and the sign of the additive synergy is indicated over the curve over variable W . \square




Figure 2: The qualitative abstraction of the Bayesian network from Figure 1.

```

procedure Process-Observation( $Q, O, sign$ ):
  for all  $A_i \in V(G)$  in  $Q$ 
  do  $sign[A_i] \leftarrow '0'$ ;
  Propagate-Sign( $Q, \emptyset, O, sign$ ).

procedure Propagate-Sign( $Q, trail, to, message$ ):
   $sign[to] \leftarrow sign[to] \oplus message$ ;
   $trail \leftarrow trail \cup \{to\}$ ;
  for each relevant neighbour  $A_i$  of  $to$  in  $Q$ 
  do  $linksign \leftarrow$  sign of influence between  $to$  and  $A_i$ ;
      $message \leftarrow sign[to] \otimes linksign$ ;
     if  $A_i \notin trail$  and  $sign[A_i] \neq sign[A_i] \oplus message$ 
     then Propagate-Sign( $Q, trail, A_i, message$ ).

```

Figure 3: The basic sign-propagation algorithm.

We would like to note that, although in the previous example the signs of the qualitative relationships were computed from the conditional probability distributions of the corresponding quantitative network, in realistic applications these signs would be elicited directly from domain experts.

2.2 Qualitative Probabilistic Reasoning

For probabilistic reasoning with a qualitative network, an efficient algorithm is available [3]. This algorithm provides for computing the effect of an observation that is entered into the network, upon the probability distributions for the other variables. It is based upon the idea of propagating and combining signs, and builds upon the properties of symmetry, transitivity and composition of qualitative influences. The algorithm is summarised in pseudo-code in Figure 3.

The algorithm takes for its input a qualitative network (Q), a variable for which an observation has become available (O), and the sign of this observation ($sign$), that is, either a '+' for the observed value o_1 or a '-' for the value o_2 . The algorithm now traces the effect of the observation throughout the network, by passing messages between neighbouring variables. For each variable A , it determines a *node sign* ' $sign[A]$ ' that indicates the direction of change in probability distribution that is occasioned by the observation; initially all node signs are set to '0'. The actual inference is started by the observed variable receiving a message with the sign of the observation. Each variable that receives a message, updates its node sign using the \oplus -operator and subsequently sends a message to each relevant neighbour; a neighbour is relevant if it is not blocked from the observed variable along the trail that is currently being followed. The sign of the message that the variable sends to a neighbour, is the \otimes -product of its own node sign and the sign of the influence that the message will traverse. This process of message passing between neighbouring variables is repeated iteratively. A trail of messages ends as soon as there are no more relevant neighbours to visit or as soon as the last message does not change the node sign of the visited variable. Since the node sign of each variable can change at most twice, once from '0' to '+', '-' or '??', and then only to '??', the process visits each variable at most twice and is guaranteed to halt in polynomial time.

The sign-propagation algorithm reviewed above serves to compute the effect of a *single* observation on the distributions for *all* other variables in a qualitative network. In realistic applications, often the effect of *multiple* simultaneous observations on a single variable is of interest. This joint effect can be computed as the \oplus -sum of the effects of the separate observations on the variable of interest [8].

Example 3 We consider the qualitative network from Figure 4. Suppose that we are interested in the effect of observing the value e_2 for the variable E on the probability distributions for the other variables in the network. Prior to the inference, the node signs of all variables are set to

'0'. Inference is started by entering the observation into the network, that is, by sending the message '- ' to the variable E . E updates its node sign to $0 \oplus - = -$ and subsequently sends the message $- \otimes + = -$ to its neighbour B . Upon receiving this message, B updates its node sign to $0 \oplus - = -$. It subsequently sends the messages $- \otimes - = +$ to C , $- \otimes ? = ?$ to A and $- \otimes + = -$ to F . Upon receiving these messages, variables C , A and F update their node signs to $0 \oplus + = +$, $0 \oplus ? = ?$ and $0 \oplus - = -$, respectively. Variable C sends the message $+ \otimes + = +$ to D . The inference ends after D has updated its node sign to $0 \oplus + = +$. The resulting node signs are indicated in the figure.

Now suppose that we are interested in the joint effect of the simultaneous observations $D = d_1$ and $E = e_2$ on the variable A . The effect of the observation $E = e_2$ on the probability distribution for A is ambiguous, as illustrated above. The effect of the observation $D = d_1$ on the distribution for A is also computed for the initial state of the network. The inference runs comparably, except that the variable A is blocked from the observed variable D : upon receiving its message from C , variable B does not send any information to A . Propagation of the observation $D = d_1$ thus results in the node sign '0' for variable A . We conclude that the combined effect of the two observations on A is $? \oplus 0 = ?$. \square

3 Situational Signs

The presence of influences with ambiguous signs in a qualitative probabilistic network is likely to give rise to ambiguous, and therefore uninformative, results upon inference, as was illustrated in the previous section. From the definitions of the \otimes - and \oplus -operators, moreover, we have that, once an ambiguous sign is encountered upon inference, it tends to spread throughout the network. The use of ambiguous signs to indicate non-monotonicity thus has undesirable consequences.

We take a closer look at the origin of the ambiguous sign of a non-monotonic influence. We observe that, for a qualitative influence of a variable A on a variable B along an arc $A \rightarrow B$ to be unambiguous, the difference $\Pr(b \mid ax) - \Pr(b \mid \bar{a}x)$ has to have the same sign for *all* combinations of values x for the set $X = \pi(B) \setminus \{A\}$ of parents of B other than A . This sign then is valid in all possible states of the network under study. More specifically, it is valid for any probability distribution over X . If the difference $\Pr(b \mid ax) - \Pr(b \mid \bar{a}x)$ yields contradictory signs for different combinations x , we have that the sign of the influence is dependent on the state of the network. The influence then is assigned the ambiguous sign '?'. In each specific state of the network, associated with a specific probability distribution $\Pr(X)$ over the combinations of values x , however, the influence of A on B is unambiguous, that is, it is either positive, negative or zero. To capture information about the current effect of a non-monotonic influence, associated with a specific state, we now introduce the concept of *situational sign*.

We consider the influence of a variable A on a variable B along an arc $A \rightarrow B$. A *positive situational sign* for the influence indicates that

- $S^?(A, B)$ and
- $[\Pr(b \mid a) - \Pr(b \mid \bar{a})]_{\Pr(X)} \geq 0$,




Figure 4: A qualitative network and its node signs after propagating the observation $E = e_2$.

where $[\Pr(b | a) - \Pr(b | \bar{a})]_{\Pr(X)}$ denotes the difference between the probabilities $\Pr(b | a)$ and $\Pr(b | \bar{a})$ in the state of the network associated with $\Pr(X)$. Negative, zero and unknown situational signs are defined analogously. We note that while the regular signs of qualitative influences and additive synergies have general validity, a situational sign is dynamic in nature: it pertains to a specific state of a network and may lose its validity as the network’s state changes. An influence with a situational sign δ now is called a *situational influence*; the sign of this situational influence is denoted ‘?’(δ)’. A qualitative probabilistic network with situational signs is termed a *situational qualitative network*.

Example 4 We consider again the Bayesian network from Figure 1 and its qualitative abstraction from Figure 2. We recall that the qualitative influence of the variable T on the variable W was found to be non-monotonic. The sign of this influence therefore depends on the state of the network. In the prior state of the network, where no evidence has been entered, we have that $\Pr(f) = 0.4$. Given this probability, we find that $\Pr(w | t) = 0.39$ and $\Pr(w | \bar{t}) = 0.51$. From the difference $\Pr(w | t) - \Pr(w | \bar{t}) = -0.12$ being negative, we conclude that, in this particular state, the influence of T on W is negative. The current sign of the situational influence therefore is ‘?’(-)’. The situational qualitative network for the prior state is shown in Figure 5.

The dynamic nature of the situational sign of the influence of T on W is demonstrated by entering the observation *true* for the variable F into the network. As a consequence of this observation, the state of the network changes. More specifically, we now have that $\Pr(f) = 1.0$. Given this probability, the difference $\Pr(w | t) - \Pr(w | \bar{t}) = 0.90 - 0.75 = 0.15$ is positive. In the new state of the network, therefore, the sign of the situational influence is ‘?’(+)'.

Once again we note that, although in the previous example the prior situational sign was computed from the corresponding quantitative network, in a realistic application it would be elicited directly from a domain expert. In the remainder of the paper, we assume that an expert has given the signs for all situational influences for the prior state of a network.

4 Inference with a Situational Qualitative Network

For inference with a regular qualitative probabilistic network, an efficient algorithm is available that is based on the idea of propagating and combining signs of qualitative influences; we reviewed this algorithm in Section 2. For inference with a situational qualitative network, we observe that the situational sign of an influence of A on B indicates the effect of a change in A ’s distribution on the probability distribution over B ’s values just like a regular sign, yet only for a particular state of the network. A situational sign can therefore be used as a regular sign upon inference *provided* that it is valid. In Section 4.1, we present a method for verifying the validity of the situational signs in a network as observations become available that cause the network to convert to another state; this method also provides for updating the signs if necessary. In Section 4.2, we incorporate this method into the sign-propagation algorithm to provide for inference with a situational qualitative network.

4.1 The Dynamics of Situational Signs

To investigate the dynamics of a situational sign, we begin by studying a network fragment of the simplest topology in which a non-monotonic influence occurs; this fragment is composed of a




Figure 5: The network from Figure 2, with the situational influence of T on W in the prior state of the network.

single variable with two mutually independent parents. For this network fragment, we show how the validity of the situational sign involved can be verified upon inference. We then extend the main idea to arrive at a method for verifying situational signs in networks of general topology and for updating them if necessary.

We consider the network fragment from Figure 6. We assume for now that the variables A and C remain independent as observations are being entered into the rest of the network. By conditioning on A and C , we find for the probability of b that

$$\begin{aligned} \Pr(b) &= \Pr(a) \cdot (\Pr(b | a) - \Pr(b | \bar{a})) + \Pr(b | \bar{a}) \\ &= \Pr(a) \cdot [\Pr(c) \cdot (\Pr(b | ac) - \Pr(b | a\bar{c}) - \Pr(b | \bar{a}c) + \Pr(b | \bar{a}\bar{c})) + \\ &\quad \Pr(b | a\bar{c}) - \Pr(b | \bar{a}\bar{c})] + \Pr(c) \cdot (\Pr(b | \bar{a}c) - \Pr(b | \bar{a}\bar{c})) + \Pr(b | \bar{a}\bar{c}). \end{aligned}$$

We observe that $\Pr(b)$ is a function of the probabilities of a and c . For a fixed $\Pr(c)$, moreover, the probability of b is linear in $\Pr(a)$. This linear function has the extremes $\Pr(b | a)$ and $\Pr(b | \bar{a})$ for $\Pr(a) = 1$ and $\Pr(a) = 0$, respectively. The sign of the difference $\Pr(b | a) - \Pr(b | \bar{a})$ now is the situational sign of the influence of A on B in the state of the network that is associated with that particular $\Pr(c)$. The sign of the difference also is the sign of the gradient of the function. We thus have that the sign of the gradient of the function that expresses $\Pr(b)$ in terms of $\Pr(a)$ at a particular, fixed, $\Pr(c)$, matches the situational sign of the influence of A on B in the associated state of the network.

In essence, there are two different manifestations of the non-monotonic influence of A on B : either the situational influence is negative for lower values of $\Pr(c)$ and positive for higher values of $\Pr(c)$, or vice versa. An example of the former manifestation is shown in Figure 7, while Figure 8 depicts an example of the latter manifestation. We observe that the manifestation from Figure 7 has associated a positive additive synergy of A and C on B : for this manifestation, we have $\Pr(b | ac) - \Pr(b | \bar{a}c) > 0$ and $\Pr(b | a\bar{c}) - \Pr(b | \bar{a}\bar{c}) < 0$ from which we find $\Pr(b | ac) + \Pr(b | \bar{a}\bar{c}) > \Pr(b | a\bar{c}) + \Pr(b | \bar{a}c)$. Similarly, the manifestation from Figure 8 corresponds with a negative additive synergy of A and C on B .

As observations are being entered into the network, the probability of c may change. When $\Pr(c)$ changes, we find another linear relationship between $\Pr(a)$ and $\Pr(b)$, with a different gradient, possibly with a different sign. If the probability of c changes, therefore, the current situational sign of the influence of A on B may become invalid. Whether or not it is invalidated, depends on the manifestation of the non-monotonic influence and on the direction of change of the probability of c . In the graph depicted in Figure 7, for example, the situational sign will definitely persist if it is negative and the probability of c decreases, or if it is positive and the probability of c increases. The reverse holds for Figure 8.

The previous observations suggest that a method for verifying whether or not a situational sign retains its validity, has to distinguish between the two possible manifestations of the underlying non-monotonic influence. We recall that these manifestations are characterised by different signs for the additive synergy involved. Now, after a change in the probability of c , we have to reconsider the difference $[\Pr(b | a) - \Pr(b | \bar{a})]_{\Pr(C)}$:

$$\begin{aligned} [\Pr(b | a) - \Pr(b | \bar{a})]_{\Pr(C)} &= \\ &\Pr(c) \cdot (\Pr(b | ac) - \Pr(b | a\bar{c}) - \Pr(b | \bar{a}c) + \Pr(b | \bar{a}\bar{c})) + \Pr(b | a\bar{c}) - \Pr(b | \bar{a}\bar{c}). \end{aligned}$$

We observe that the difference $[\Pr(b | a) - \Pr(b | \bar{a})]_{\Pr(C)}$ is a linear function in $\Pr(c)$. We further observe that the sign of the gradient of the function equals the sign of the additive synergy of




Figure 6: A fragment of a situational network, consisting of variable B and its parents A and C , with $S^{?(\delta_1)}(A, B)$ and $Y^{\delta_2}(\{A, C\}, B)$.




Figure 7: An example $\Pr(b)$ as a function of $\Pr(a)$ and $\Pr(c)$, with $S^?(A, B), S^+(C, B)$ and $Y^+(\{A, C\}, B)$.




Figure 8: An example $\Pr(b)$ as a function of $\Pr(a)$ and $\Pr(c)$, with $S^?(A, B), S^+(C, B)$ and $Y^-(\{A, C\}, B)$.

A and C on B . Now suppose that $Y^+(\{A, C\}, B)$. The gradient of the function then is positive and the manifestation exemplified in Figure 7 holds. If the probability of c increases as a result of newly entered observations, a positive situational sign will definitely remain valid. If, on the other hand, $Y^-(\{A, C\}, B)$ and the probability of c increases, then a negative situational sign will retain its validity. We thus have that, upon an increase of $\Pr(c)$, a situational sign δ_1 persists if $\delta_1 = + \otimes \delta_2$, where δ_2 is the sign of the additive synergy involved; otherwise, the situational sign becomes unknown and δ_1 should be changed to '?. Similar observations hold for a decreasing probability of c . We conclude that the updating of a situational sign is captured by

$$\delta_1 \leftarrow \delta_1 \oplus (\text{sign}[C] \otimes \delta_2).$$

Without further substantiation, we extend the previous observations to the more general situation in which B has two or more mutually independent parents. We consider a variable B with parents A and C_i , $i = 1, \dots, n, n \geq 1$, with $S^{?(\delta)}(A, B)$ and $Y^{\delta_i}(\{A, C_i\}, B)$. Then, updating the situational sign of the influence of A on B is captured by

$$\delta \leftarrow \delta \oplus_{i=1}^n (\text{sign}[C_i] \otimes \delta_i).$$

In our analysis so far, we assumed that the two parents A and C of the variable B are mutually independent and remain to be so as evidence is entered into the network. In general, however, A and C can be (conditionally) dependent. Variable A then not only influences B directly, but also indirectly through C . The situational influence of A on B , however, pertains to the direct influence in isolation even though a change in the probability of c may affect its sign. When a change in the probability of a causes a change in the probability of c which in turn influences the probability of b , the indirect influence on B is processed separately and the sign of the net influence equals the composition of the signs of the two influences.

4.2 The Adapted Sign-Propagation Algorithm

The basic sign-propagation algorithm for inference with a qualitative network has to be adapted to render it applicable to situational qualitative networks. In essence, the following modifications are required. First, of non-monotonic influences, the situational signs should be used for the propagation instead of the original '?. Due to the process of sign propagation, moreover, it may occur that a sign is propagated from a variable A to a variable B , while the fact that the probability distribution of another parent of B has changed does not become apparent until later in the inference. It may then turn out that the situational sign should have been updated and that incorrect signs were propagated over the influence with the situational sign. The algorithm therefore has to verify the validity of a situational sign as soon as information to this end becomes available and, if the situational sign is updated, to restart the inference with the updated network. Since a situational sign can change at most twice, from '0' to '+' or '-' and then to '?', the number of restarts is limited. The adapted algorithm is summarised in pseudo-code in Figure 9.

The algorithm takes for its input a situational qualitative network (Q), a variable for which an observation has become available (O), and the sign of this observation ($sign$). The algorithm constructs from the network the set ARC_{nm} of all arcs with an associated non-monotonic influence, and the set COP_{nm} of all nodes that are co-parents of a variable exerting a non-monotonic influence. The function $COP-ARC_{nm}(A)$ takes for its argument a variable and returns all arcs with a non-monotonic influence exerted by a co-parent of this variable. While the procedure 'Process-Observation' is identical to the same procedure in the regular algorithm, the procedure 'Propagate-Sign' is modified. After 'sign[to] \leftarrow sign[to] \oplus message', which may have led to a change of node sign, a call to the new procedure 'Determine-Effect-On' is added. This procedure updates the situational signs of the network and, if necessary, restarts the inference.

Like the basic sign-propagation algorithm, the adapted algorithm serves to compute the effect of a single observation on the distributions for all other variables in a network. The joint effect, on a variable of interest, of multiple simultaneous observations can again be computed as the \oplus -sum of the effects of the separate observations. However, when a situational sign changes during

$$\begin{aligned}
ARC_{nm} &= \{A_i \rightarrow A_j \mid S^{2(\delta)}(A_i, A_j)\} \\
COP_{nm} &= \{A_k \mid A_k \in \pi(A_j) \setminus \{A_i\}, A_i \rightarrow A_j \in ARC_{nm}\} \\
COP-ARC_{nm}(A) &= \{A_i \rightarrow A_j \mid A_i \rightarrow A_j \in ARC_{nm}, A_j \in \sigma(A), A_i \neq A\};
\end{aligned}$$

procedure Process-Observation($Q, O, sign$):

for all $A_i \in V(G)$ in Q
do $sign[A_i] \leftarrow '0'$;
Propagate-Sign($Q, \emptyset, O, sign$).

procedure Propagate-Sign($Q, trail, to, message$):

$sign[to] \leftarrow sign[to] \oplus message$;
 $trail \leftarrow trail \cup \{to\}$;
Determine-Effect-On(Q, to);
for each relevant neighbour A_i of to in Q
do $linksign \leftarrow$ (situational) sign of influence between to and A_i ;
 $message \leftarrow sign[to] \otimes linksign$;
if $A_i \notin trail$ and $sign[A_i] \neq sign[A_i] \oplus message$
then Propagate-Sign($Q, trail, A_i, message$).

procedure Determine-Effect-On(Q, A_i):

if $A_i \in COP_{nm}$
then for all $A_j \rightarrow A_k \in COP-ARC_{nm}(A_i)$
do Verify-Update($S^{2(\delta)}(A_j, A_k)$);
if a δ changes
then $Q \leftarrow Q$ with adapted signs;
return Process-Observation($Q, O, sign$).

Figure 9: The adapted sign-propagation algorithm.

the propagation of one of the observations, the propagation of all other observations has to be performed anew with the adapted network before establishing the joint effect. Again, because a situational sign can change at most twice, the number of restarts is limited.

Example 5 We consider the situational qualitative network from Figure 10; the network is identical to the regular qualitative network from Figure 4, except that it is supplemented with a situational sign for the non-monotonic influence of variable A on variable B for the prior state of the network. From the situational network, the sets $ARC_{nm} = \{A \rightarrow B\}$ and $COP_{nm} = \{C\}$ are established. Suppose that we are again interested in the effect of observing the value e_2 for the variable E on the probability distributions for the other variables in the network. Inference is started by sending the message ‘-’ to the variable E . E updates its node sign to $0 \oplus - = -$ and subsequently sends the message $- \otimes + = -$ to its neighbour B . Variable B updates its node sign to $0 \oplus - = -$ and subsequently sends the messages $- \otimes - = +$ to C , $- \otimes + = -$ to A , and $- \otimes + = -$ to F . Upon receiving these messages, variables C , A and F update their node signs to $0 \oplus + = +$, $0 \oplus - = -$ and $0 \oplus - = -$, respectively. The algorithm now establishes that C is a co-parent, with A , of B and that the influence of A on B is non-monotonic. Because the node sign of C has changed, the validity of the situational sign of the influence of A on B needs to be verified. The algorithm therefore checks if the current situational sign equals the product of the node sign of C and the sign of the additive synergy involved. Since $+ = + \otimes +$, the algorithm concludes that the situational sign remains valid. The inference resumes with variable C sending the message $+ \otimes + = +$ to D . D updates its node sign to $0 \oplus + = +$ and the inference ends. The resulting node signs are indicated in the figure. \square

We conclude from Examples 3 and 5 that inference with a situational network can yield more informative results than inference with the corresponding regular qualitative network.




Figure 10: An example situational qualitative network.

5 An Experimental Study

In the previous section we demonstrated that a situational network can yield more informative results upon inference than a corresponding regular qualitative network. The example that we used, however, pertained to a small, artificially constructed network. In this section, we investigate the practicability of situational signs by studying the effects of their introduction into a real-life qualitative network in the field of oesophageal cancer. In Section 5.1 we provide some background information on this network. In Section 5.2 we describe the performance of the network before and after the introduction of situational signs, for a number of real patients. In Section 5.3 we review the results of our study.

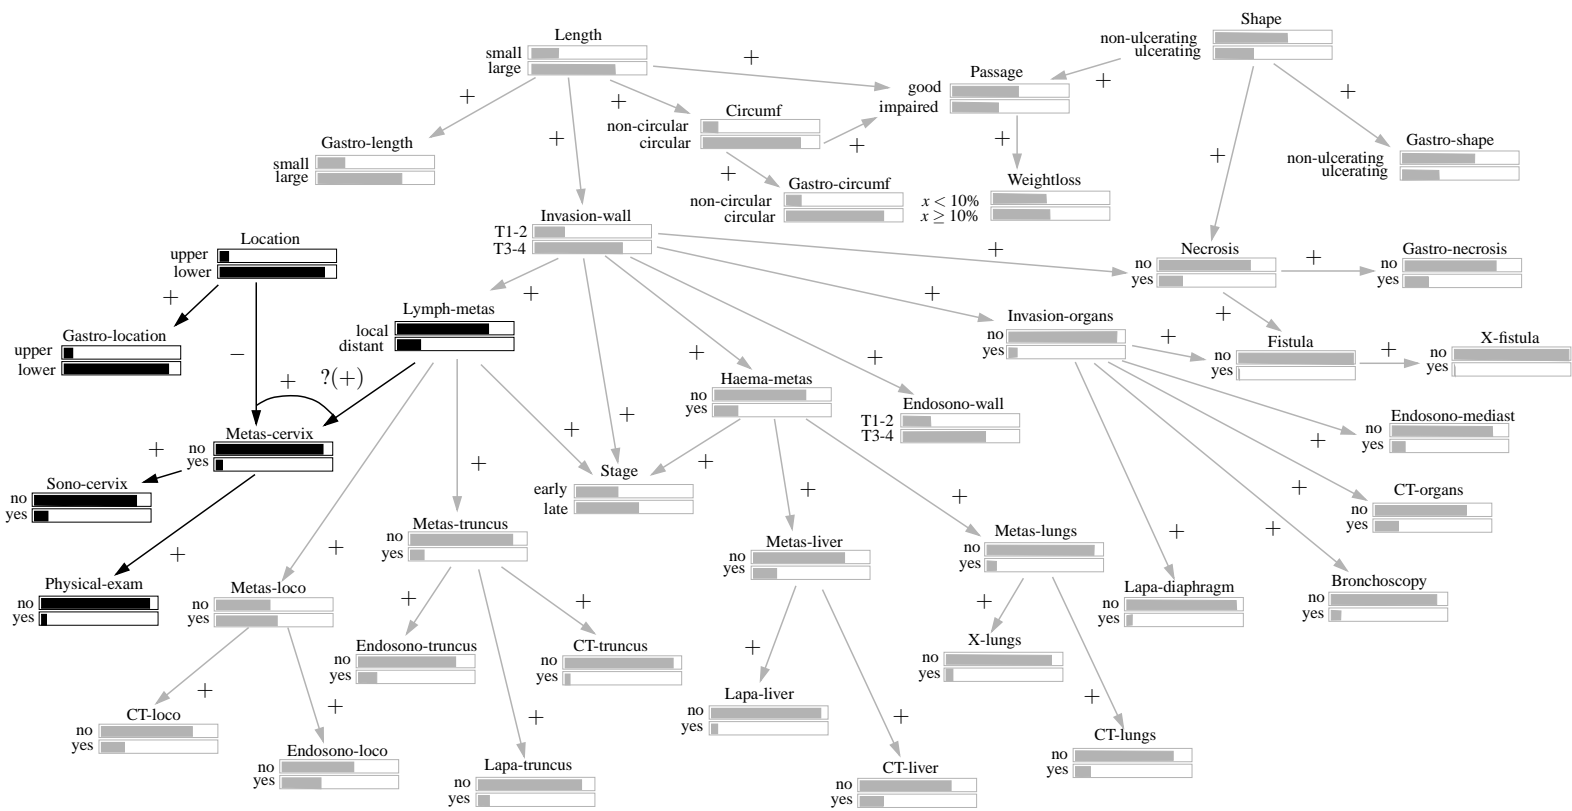
5.1 The Oesophageal Cancer Network

A chronic lesion of the inner wall of the oesophagus may develop into a malignant tumour. This tumour invades the oesophageal wall and upon further growth may invade adjacent organs. In time, the tumour may give rise to metastases in lymph nodes and to secondary tumours in the lungs and the liver. The depth of invasion and extent of metastasis indicate how far the cancer has progressed or, phrased alternatively, in which stage it is. To establish the stage of a patient’s cancer, various diagnostic tests are performed, ranging from multiple biopsies of the primary tumour to a radiograph of the patient’s lungs. The state-of-the-art knowledge about oesophageal cancer was captured in a Bayesian network [9]. This network includes 42 statistical variables and some thousand conditional probabilities. Its main diagnostic variable is the variable *Stage*, classifying a patient’s cancer in one of six possible stages of disease. The leaves of the network capture the possible results of the different diagnostic tests.

For our study, we abstracted the oesophageal cancer network to a qualitative network. To this end, we first summarised all variables into binary variables, building upon our knowledge of the domain. The original six-valued variable *Stage*, for example, was translated into the binary variable *Stage* with the values *early* and *late*. We further defined orderings on the values of the resulting binary variables. We took, for example, $early < late$. Given these orderings, we established the signs for the influences and the additive synergies between the variables from the probabilities specified for the original network. We decided to delete the arcs that were associated with a (nearly) zero influence, which resulted in the removal of 15 arcs and two nodes. Figure 11 shows the binary quantitative oesophageal cancer network as well as its qualitative abstraction. For each variable, its name, its values, and its prior probability distribution are shown; for each arc, moreover, the sign of the associated qualitative influence is depicted.

The qualitative oesophageal cancer network includes a single non-monotonic influence, between the variables *Lymph-metas* and *Metas-cervix*. The variable *Lymph-metas* models whether or not the primary tumour has metastasized to distant lymph nodes; the variable *Metas-cervix* models whether or not the lymph nodes in the neck are affected by the cancer. The actual sign of the influence between the two variables depends on the value of the variable *Location*. This variable models whether the primary tumour resides in the upper one-third of the oesophagus, or in the lower two-third. The lymph nodes in the neck are considered local for a primary tumour in the upper one-third of the oesophagus, and distant otherwise. For a tumour located in the upper

Figure 11: The combined presentation of the binary and the qualitative oesophagus network.



one-third of the oesophagus, the presence of metastases in distant lymph nodes has a negative effect on the probability of metastases in the neck; if, on the other hand, the primary tumour is located in the lower two-third of the oesophagus, the presence of distant lymphatic metastases has a positive effect on this probability. In the initial state of the network, where no evidence has been entered, the probability of the tumour being located in the lower two-third of the oesophagus is quite high, and the situational sign of the non-monotonic influence, accordingly, is '+'.

The non-monotonic influence between the variables *Lymph-metas* and *Metas-cervix* resides in a pivotal position in the qualitative oesophageal cancer network. For establishing its stage, knowledge of the extent of the lymphatic metastases of a patient's cancer is of primary importance. The presence or absence of metastases in the neck and their classification as local or distant, therefore, play a crucial role in the staging. The non-monotonic influence now links the variable *Lymph-metas* to the part of the network pertaining to metastases in the neck.

The variables *Physical-exam* and *Sono-cervix* model the diagnostic tests that are generally performed to establish the presence or absence of lymphatic metastases in the neck; they represent the findings from a physical examination of the neck and from a sonography of the neck, respectively. The location of the primary tumour is established through a gastroscopic examination of the oesophagus. The result of this examination is captured by the variable *Gastro-location*. The variables *Physical-exam* and *Sono-cervix* upon observation influence the node sign of *Lymph-metas*. A simultaneous observation for the variable *Gastro-location* bears no influence on *Lymph-metas*, because, in the prior state of the network, *Gastro-location* is independent of *Lymph-metas*. The node sign of the variable *Location* is influenced by observations for all three variables. We note that the node sign of this variable is instrumental in updating the situational sign of the non-monotonic influence between *Lymph-metas* and *Metas-cervix* after observations have caused the network's state to change.

5.2 The Effect of the Introduction of Situational Signs

To gain insight into the practicability of situational signs, we study the performance of the qualitative oesophageal cancer network, before and after the introduction of a situational sign for the influence between *Lymph-metas* and *Metas-cervix*. In doing so, we focus on the part of the network that serves for interpreting the findings with regard to metastases in the neck; the part of the network under study is indicated in black in Figure 11. We investigate whether useful information from this part of the network is propagated towards the variable *Lymph-metas* upon inference. In our study, we use the data of 156 real patients diagnosed with cancer of the oesophagus. We first demonstrate, as an example, the effect of introducing the situational sign for a single patient, after which we summarise the effect for all patients from our data collection.

Example 6 For patient 90-1042, gastroscopic examination showed a primary tumour in the lower two-third of the oesophagus. Physical examination did not reveal any enlarged lymph nodes in the patient's neck. A sonography was not performed. The two available observations are entered into the network as a '+' for the variable *Gastro-location* and a '-' for *Physical-exam*, respectively. Upon inference with the regular qualitative network without the situational sign, the variable *Lymph-metas* receives the message $- \otimes + \otimes ? = ?$ from *Physical-exam*. The observation of *Gastro-location* does not affect the node sign of *Lymph-metas* and inference results in an overall influence of sign '?' on this variable.

In the situational oesophageal cancer network, the influence between *Lymph-metas* and *Metas-cervix* is supplemented with a situational sign. We note that *Metas-cervix* has the variable *Location* for its other parent. Because the two available observations upon inference change the node sign of *Location*, the sign-propagation algorithm identifies that the situational sign needs updating. The node sign of *Location* captures the combined effect of the two observations: since both observations have a positive effect on *Location*, its node sign is '+'. The additive synergy of *Location* and *Lymph-metas* on *Metas-cervix* also is '+'. Updating the situational sign of the influence between *Metas-cervix* and *Lymph-metas* now gives $+ \oplus (+ \otimes +) = +$, that is, the situational sign retains

its validity and, hence, its informativeness. The part of the network that pertains to metastases in the neck now exerts an overall influence of sign $- \otimes + \otimes + = -$ on the variable *Lymph-metas*.

Note that, if the node sign of *Location* would have changed to '- ', then the situational sign would have been updated to '?'. The observation for the variable *Physical-exam* would then have exerted an ambiguous influence on *Lymph-metas*. A similar observation holds if the node sign of *Location* would have changed to '?'. Such a change occurs if the available observations exert discordant influences on *Location*, for example *Physical-exam = yes* and *Gastro-location = lower*. \square

The data collection that we have available for our study includes the medical records of 156 patients diagnosed with oesophageal cancer. For 11 of these patients we have that either *Sono-cervix = yes* and *Physical-exam = yes* or that one of these observations is *yes* and the other one is unknown. In the sequel we will call such combinations of observations consistently positive; negative consistency has an analogous meaning. For 4 of these patients we further have that *Gastro-location = upper* and for 7 of them we have that *Gastro-location = lower*. For another 7 patients we have that *Sono-cervix* and *Physical-exam* are consistently negative, and *Gastro-location = upper*. For 52 patients, we have that the observations for *Sono-cervix* and *Physical-exam* are consistently negative, and *Gastro-location = lower*. For one patient contradictory results were found from the sonography and the physical examination. For the remaining 85 patients, no observations are available from a sonography of the neck or from a physical examination. For 2 of these patients we have that *Gastro-location = lower*; for the other 83 patients we have that *Gastro-location = upper*. These statistics are summarised in Table 2.

For the 85 (55%) patients for whom no observations are available for *Sono-cervix* and *Physical-exam*, the part of the network under study does not partake in establishing the node sign of *Lymph-metas*. The non-monotonic influence between *Lymph-metas* and *Metas-cervix*, therefore, is not used upon inference for these patients. For the remaining 71 (45%) patients, inference with the regular qualitative oesophageal cancer network results in an unknown influence on the variable *Lymph-metas*.

We now address inference with the corresponding situational network. For the 85 patients without any observations for *Sono-cervix* and *Physical-exam*, the availability of the situational sign makes no difference. For the other 71 patients, the situational sign of the non-monotonic influence is used upon inference, instead of the original '?'. For all these patients, the available observations result in a change of the node sign of the variable *Location*, thereby enforcing the situational sign to be updated. For 19 (12% of all patients) of the 71 patients for whom at least one observation is available for the variables *Sono-cervix* and *Physical-exam*, the node sign of *Location* changes to a '- ' or a '?'. As for these patients the situational sign is updated to '?', inference results in an unknown effect on the variable *Lymph-metas*. For the remaining 52 (33%) patients, the node sign of *Location* changes to a '+ ' and the situational sign retains its validity. For these patients, inference yields an overall negative influence on the variable *Lymph-metas* and, hence, an informative result. The inference results obtained with the regular and situational qualitative oesophageal cancer networks are summarised in Table 3.

Table 2: The available observations for the relevant variables for 156 patients.

<i>Sono-cervix</i> and <i>Physical-exam</i>	<i>Gastro-location</i>	
	<i>upper</i>	<i>lower</i>
consistently positive	4	7
consistently negative	7	52
inconsistent	-	1
not observed	2	83

Table 3: The signs propagated from the part of the network under consideration to the variable *Lymph-metas*.

	+	-	?	0
regular	-	-	71 (45%)	85 (55%)
situational	-	52 (33%)	19 (12%)	85 (55%)

5.3 Discussion

Before the introduction of situational signs into the qualitative oesophageal cancer network, for 45% of the patients ambiguous information was propagated from the part of the network under consideration. This percentage reduced to 12% after introducing a situational sign for an important non-monotonic influence in the network. We may conclude that the introduction of situational signs served to considerably increase the expressive power of the qualitative oesophageal cancer network.

In our experiment, we observed that, for all 71 patients for whom one or more observations are available for the variables *Sono-cervix* and *Physical-exam*, the situational sign had to be updated upon inference. For 52 of these patients, the sign proved to retain its validity. Also for the other 85 patients, the situational sign was updated upon inference, even though it was not used for further propagation. For two of these patients, the situational sign changed to '?' and for 83 of these patients, the situational sign remained a '+'. We thus find that for a total of 135 (87%) patients, the situational sign retained its validity after updating. This apparent robustness is not coincidental. The situational sign depends on the prior probability of the tumour being located in the lower two-third of the oesophagus. Given the positive additive synergy of the variables *Location* and *Lymph-metas* on *Metas-cervix*, this probability, being quite high, causes the prior situational sign to be positive. This prior probability, moreover, is quite high because we are more likely to find observations that lead to a change of the node sign of the variable *Location* to '+'; observations like this are exactly those that do not induce a change of the situational sign.

6 Conclusions

Qualitative probabilistic networks capture the probabilistic influences among their variables by means of qualitative signs. If an influence between two variables is non-monotonic, it has associated the ambiguous sign '?', even though the effect of the influence is unambiguous in any particular state of the network. The presence of such ambiguous signs tends to lead to ambiguous and, hence, uninformative results upon inference. In this paper, we introduced the concept of situational sign to capture information about the current effect of non-monotonic influences. We showed that situational signs can be used upon inference and may effectively forestall ambiguous results. We identified conditions under which situational signs retain their validity and presented a method for updating them if necessary. Although we studied the dynamics of situational signs in networks where the non-monotonicity involved originates from a single variable, the presented ideas and methods are readily generalised to networks where the non-monotonicity is provoked by more than one variable.

To investigate the practicability of situational signs, we studied the effect of their introduction into a real-life qualitative network in the field of oncology. In this study, we compared the performance of the network before and after the introduction of situational signs, using the data from 156 patients. We found that the introduction of situational signs served to considerably increase the expressive power of the network under study. As our network is in no aspect exceptional, we expect similar results for other real-life qualitative networks in a variety of problem domains.

Acknowledgment

This research was supported by the Netherlands Organisation for Scientific Research (NWO).

References

- [1] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Palo Alto, 1988.
- [2] M.P. Wellman. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, vol. 44, pp. 257–303, 1990.
- [3] M.J. Druzdzel, M. Henrion. Efficient reasoning in qualitative probabilistic networks. *Proceedings of the Eleventh National Conference on Artificial Intelligence*. AAAI Press, Menlo Park, California, pp. 548–553, 1993.
- [4] S. Renooij, L.C. van der Gaag. From qualitative to quantitative probabilistic networks. In: A. Darwiche, N. Friedman, editors, *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, San Francisco, pp. 422–429, 2002.
- [5] M.J. Druzdzel, L.C. van der Gaag. Elicitation of probabilities for belief networks: combining qualitative and quantitative information. In: Ph. Besnard, S. Hank, editors, *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, San Francisco, pp. 141–148, 1995.
- [6] S. Renooij, L.C. van der Gaag. Enhancing QPNs for trade-off resolution. In: K. Laskey, H. Prade, editors, *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, San Francisco, pp. 480–487, 1999.
- [7] M.J. Druzdzel, M. Henrion. Intercausal reasoning with uninstantiated ancestor nodes. In: D. Heckerman, A. Mamdani, editors, *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, San Mateo, California, pp. 317–325, 1993.
- [8] M.J. Druzdzel. *Probabilistic Reasoning in Decision Support Systems: From Computation to Common Sense*. PhD thesis, Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, Pennsylvania, 1993.
- [9] L.C. van der Gaag, S. Renooij, C.L.M. Witteman, B.M.P. Aleman, B.G. Taal. Probabilities for a probabilistic network: a case study in oesophageal cancer. *Artificial Intelligence in Medicine*, vol. 25, pp. 123–148, 2002.