Introducing Situational Influences in QPNs

Janneke H. Bolt, Linda C. van der Gaag, and 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 models the probabilistic relationships between its variables by means of signs. Non-monotonic influences are modelled by the ambiguous sign '?', which indicates that the actual sign of the influence depends on the current state of the network. The presence of influences with such ambiguous signs tends to lead to ambiguous results upon inference. In this paper we introduce the concept of situational influence into qualitative networks. A situational influence is a non-monotonic influence supplemented with a sign that indicates its effect in the current state of the network. We show that reasoning with such situational influences may forestall ambiguous results upon inference; we further show how these influences change as the current state of the network changes.

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. The usually large number of probabilities required for a Bayesian network, however, tends to pose a major obstacle to the construction [2]. *Qualitative probabilistic networks* (QPNs), introduced as qualitative abstractions of Bayesian networks [3], do not suffer from this quantification obstacle. Like a Bayesian network, a qualitative network encodes variables and the probabilistic relationships between them in a directed graph; the relationships between the variables are not quantified by conditional probabilities as in a Bayesian network, however, but are summarised by qualitative signs instead. For inference with a qualitative probabilistic network an efficient algorithm is available, based on the idea of propagating and combining signs [4].

Although qualitative probabilistic networks do not suffer from the obstacle of requiring a large number of probabilities, their high level of abstraction causes some lack of representation detail. As a consequence, for example, qualitative networks do not provide for modelling *non-monotonic* influences in an informative way. 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. Such a non-monotonic influence is modelled by the ambiguous sign '?'. The presence of influences with such ambiguous signs typically leads to ambiguous, and thereby uninformative, results upon inference.

T.D. Nielsen and N.L. Zhang (Eds.): ECSQARU 2003, LNAI 2711, pp. 113–124, 2003.

Non-monotonicity of an influence in essence indicates that the influence cannot be captured by an unambiguous sign of general validity. In each particular state of the network, however, the influence is unambiguous. In this paper we extend the framework of qualitative probabilistic networks with *situational influences* that capture information about the current effect of non-monotonic influences. We show that these situational influences can be used upon inference and may effectively forestall ambiguous results. Because the sign of a situational influence depends on the current state of the network, we investigate how it changes as the state of the network changes. We then adapt the standard propagation algorithm to inference with networks including situational influences.

The remainder of this paper is organised as follows. Section 2 provides some preliminaries on qualitative probabilistic networks. Section 3 introduces the concept of situational influence. Its dynamics are described in Sect. 4, which also gives an adapted propagation algorithm. The paper ends with some conclusions and directions for further research in Sect. 5.

2 Preliminaries

Qualitative probabilistic networks were introduced as qualitative abstractions of Bayesian networks. Before reviewing qualitative networks, therefore, we briefly address their quantitative counterparts.

A Bayesian network is a concise representation of a joint probability distribution Pr on a set of statistical variables. In the sequel, (sets of) variables are denoted by upper-case letters. For ease of exposition, we assume all variables to be binary, writing a for $A = true$ and $\bar{a}$ for $A = false$. We further assume that $true > false$. Each variable is now represented by a node in an acyclic directed graph. The probabilistic relationships between the variables are captured by the digraph's set of arcs. Associated with each variable A is a set of (conditional) probability distributions $\Pr(A \mid \pi(A))$ describing the influence of the parents $\pi(A)$ of A on the probability distribution for A itself.

Example 1. We consider the small Bayesian network shown in Fig. 1. The network represents a fragment of fictitious knowledge about the effect of training and fitness on one's feeling of well-being. Node T models whether or not one has undergone a training session, node F captures one's fitness, and node W models whether or not one has a feeling of well-being. □

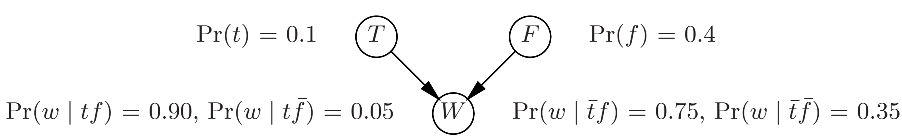

Fig. 1. An example Bayesian network, modelling the influences of training (T) and fitness (F) on a feeling of well-being (W)

In its initial state where no observations for variables have been entered, a Bayesian network captures a prior probability distribution. As such evidence becomes available, the network converts to another state and then serves to represent the posterior distribution given the evidence.

Qualitative probabilistic networks bear a strong resemblance to Bayesian networks. A qualitative network also comprises an acyclic digraph 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*, capturing features of the existing, albeit unknown, joint distribution Pr [3].

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 a node A on a node B along an arc $A \rightarrow B$, denoted $S^+(A, B)$, expresses that observing a high value for A makes the higher value for B more likely, regardless of any other direct influences on B, that is

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

for any combination of values x for the set $\pi(B) \setminus \{A\}$ of parents of B other than A. A negative qualitative influence, denoted S^-, and a zero qualitative influence, denoted S^0, are defined analogously. A non-monotonic or unknown influence of node A on node B is denoted by $S^?(A, B)$.

The set of all influences of a qualitative network exhibits various important properties [3]. 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 signs of qualitative influences along a trail without head-to-head nodes combine into a sign for the net influence with the $\otimes$-operator from Table 1. The property of *composition* asserts that the signs of multiple influences between two nodes along parallel trails combine into a sign for the net influence with the $\oplus$-operator.

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

$\otimes$	+	−	0	?
+	+	−	0	?
−	−	+	0	?
0	0	0	0	0
?	?	?	0	?

$\oplus$	+	−	0	?
+	+	?	+	?
−	?	−	−	?
0	+	−	0	?
?	?	?	?	?

A qualitative probabilistic network further includes *additive synergies*. An additive synergy expresses how two nodes interact in their influence on a third node. For example, a *positive additive synergy* of a node A and a node B on a common child C, denoted $Y^+(\{A, B\}, C)$, expresses that the joint influence of A and B on C exceeds the sum of their separate influences regardless of any other direct 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 parents of C other than A and B. 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 nodes A and B on a common child C is denoted by $Y^?(\{A, B\}, C)$.

Example 2. We consider the qualitative abstraction of the Bayesian network from Fig. 1. From the conditional probability distributions specified for node W, we have that $\Pr(w \mid tf) - \Pr(w \mid t\bar{f}) \geq 0$ and $\Pr(w \mid \bar{t}f) - \Pr(w \mid \bar{t}\bar{f}) \geq 0$, and therefore that $S^+(F, W)$: fitness favours well-being regardless of training. We further have that $\Pr(w \mid tf) - \Pr(w \mid \bar{t}f) > 0$ and $\Pr(w \mid t\bar{f}) - \Pr(w \mid \bar{t}\bar{f}) < 0$, and therefore that $S^?(T, W)$: the effect of training on well-being depends on one's fitness. From $\Pr(w \mid tf) + \Pr(w \mid \bar{t}\bar{f}) \geq \Pr(w \mid t\bar{f}) + \Pr(w \mid \bar{t}f)$, to conclude, we find that $Y^+(\{T, F\}, W)$. The resulting qualitative network is shown in Fig. 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. □

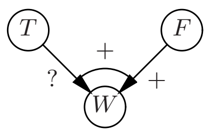

Fig. 2. The qualitative abstraction of the Bayesian network from Fig. 1

We would like to note that, although in the previous example the qualitative relationships between the variables are computed from the conditional probabilities of the corresponding quantitative network, in realistic applications these relationships are elicited directly from domain experts.

For inference with a qualitative probabilistic network, an efficient algorithm based on the idea of propagating and combining signs is available [4]. This algorithm traces the effect of observing a value for a node upon the other nodes in a network by message passing between neighbouring nodes. The algorithm is summarised in pseudo-code in Fig. 3. For each node V, a *node sign* 'sign[V]' is determined, indicating the direction of change in its probability distribution occasioned by the new observation; initial node signs equal '0'. Observations are entered as a '+' for the observed value *true*, or a '−' for the value *false*. Each node receiving a message updates its sign using the ⊕-operator and subsequently sends a message to each neighbour that is not independent of the observed node. The sign of this message is the ⊗-product of the node's (new) sign and the sign of the influence it traverses. This process of message passing between neighbours is repeated throughout the network, building on the properties of symmetry, transitivity, and composition of influences. Since each node can change its sign at most twice (once from '0' to '+', '−' or '?', and then only to '?'), the process visits each node at most twice and therefore halts in polynomial time.

Example 3. We consider the qualitative network shown in Fig. 4. Suppose that we are interested in the effect of observing the value *false* for node A upon the other nodes in the network. Prior to the inference, the node signs for all nodes

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

procedure Propagate-Sign(Q,$trail$,to,$message$):
  sign[to] ← sign[to] ⊕ $message$;
  $trail$ ← $trail$ ∪ {to};
  for each neighbour V_i of to in Q
  do $linksign$ ← sign of influence between to and V_i;
    $message$ ← sign[to] ⊗ $linksign$;
    if $V_i \notin trail$ and sign[V_i] ≠ sign[V_i] ⊕ $message$
    then Propagate-Sign(Q,$trail$,V_i,$message$).

Fig. 3. The sign-propagation algorithm

are set to '0'. Inference now starts with node A receiving the message '$-$'. Node A updates its node sign to $0 \oplus - = -$, and subsequently computes the messages to be sent to its neighbours E, B and D. To node E, node A sends the message $- \otimes - = +$. Upon receiving this message, node E updates its node sign to $0 \oplus + = +$. Node E does not propagate the message it has received from A to node B because A and B are independent on the trail $A \rightarrow E \leftarrow B$. To node B, node A sends the message $- \otimes ? = ?$. Upon receiving this message, node B updates its node sign to $0 \oplus ? = ?$. Node B subsequently computes the message $? \otimes + = ?$ for E. Upon receiving this message, node E updates its node sign to $+ \oplus ? = ?$. Node B does not propagate the message it has received from A to node C because A and C are independent on the trail $A \rightarrow B \leftarrow C$. Exploiting the property of symmetry, node A sends the message $- \otimes + = -$ to node D. Upon receiving this message, node D updates its node sign to $0 \oplus - = -$. Node D subsequently computes the message $- \otimes + = -$ for C. Upon receiving this message, node C updates its node sign to $0 \oplus - = -$. Node C then sends the message $- \otimes - = +$ to B, upon which node B should update its node sign to $? \oplus + = ?$. Since this update would not change the node sign of B, the propagation of messages halts. The node signs resulting from the inference are shown in the network's nodes in Fig. 4. □

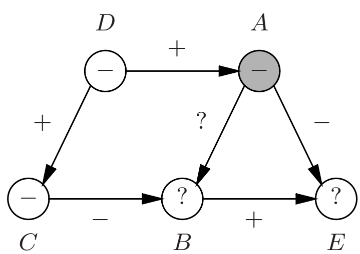

Fig. 4. A qualitative network and its node signs after the observation $A = false$

3 Situational Influences

The presence of influences with ambiguous signs in a qualitative network is likely to give rise to uninformative ambiguous results upon inference, as illustrated in Example 3. We take a closer look at the origin of these ambiguous signs. We observe that a qualitative influence of a node A on a node B along an arc $A \rightarrow B$ is only unambiguous if the difference $\Pr(b \mid ax) - \Pr(b \mid \bar{a}x)$ has the same sign for *all* combinations of values x for the set $X = \pi(B) \setminus \{A\}$. As soon as the difference $\Pr(b \mid ax) - \Pr(b \mid \bar{a}x)$ yields contradictory signs for different combinations x, the influence is non-monotonic and is assigned the ambiguous sign '?'. In each specific state of the network, associated with a specific probability distribution $\Pr(X)$ over all combinations x, however, the influence of A on B is unambiguous, that is, either positive, negative or zero. To capture the current sign of a non-monotonic influence in a specific state, we introduce the concept of *situational influence* into the formalism of qualitative probabilistic networks.

We consider a qualitative network as before and consider the evidence e entered so far. A *positive situational influence* of a node A on a node B given e, denoted $S_e^{?(+)}(A, B)$, is a non-monotonic influence of A on B for which

$$\Pr(b \mid ae) - \Pr(b \mid \bar{a}e) \geq 0\,.$$

In the sequel we omit the subscript e from $S_e^{?(+)}$ as long as ambiguity cannot occur. A negative situational influence, denoted $S^{?(-)}$, and a zero situational influence, denoted $S^{?(0)}$, are defined analogously. An unknown situational influence of node A on node B is denoted by $S^{?(?)}(A, B)$. The sign between the brackets will be called *the sign of the situational influence*. A qualitative network extended with situational influences will be called a *situational qualitative network*. Note that while the signs of qualitative influences and additive synergies have general validity, the signs of situational influences pertain to a specific state of the network and depend on $\Pr(X)$.

Example 4. We consider once again the network fragment from Fig. 1 and its qualitative abstraction shown in Fig. 2. The qualitative influence of node T on node W was found to be non-monotonic. Its sign 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 $\Pr(w \mid t) = 0.39$ and $\Pr(w \mid \bar{t}) = 0.51$. From the difference $\Pr(w \mid t) - \Pr(w \mid \bar{t}) = -0.12$ being negative, we conclude that the influence of node T on node W is negative in this particular state. The current sign of the influence is therefore '$-$'. The situational qualitative network for the prior state is shown in Fig. 5. The dynamic nature of the sign of the situational influence is illustrated by a change from '$-$' to '$+$' after, for example, the observation $F = \textit{true}$ is entered into the network, in which case $\Pr(w \mid tf) - \Pr(w \mid \bar{t}f) = 0.90 - 0.75 = 0.15$. □

Once again we note that, although in the previous example the sign of the situational influence is computed from the quantitative network, in a realistic application it would be elicited directly from a domain expert. In the remainder

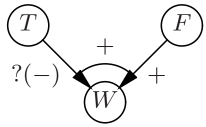

Fig. 5. The network from Fig. 2, now with the prior situational influence of T on W

of the paper, we assume that the expert has given the signs of the situational influences for the prior state of the 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, as reviewed in Sect. 2. For inference with a situational qualitative network, we observe that the sign of a situational influence indicates the sign of the original qualitative influence in the current state of the network. After an observation has been entered into a situational network, therefore, the signs of the situational influences can in essence be propagated as in regular qualitative networks, *provided* that these signs are still valid in the new state of the network. In this section we discuss how to verify the validity of the sign of a situational influence as observations become available that cause the network to convert to another state. In addition, we show how to incorporate this verification into the sign propagation algorithm.

4.1 Dynamics of the Signs of Situational Influences

We begin by investigating the simplest network fragment in which a non-monotonic qualitative influence can occur, consisting of a single node with two independent parents. We show for this fragment how the validity of the sign of the situational influence can be verified during inference by exploiting the associated additive synergy. We then extend the main idea to more general situational networks.

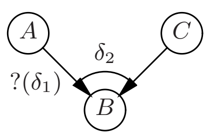

Fig. 6. A fragment of a situational network, consisting of node B and its parents A and C, with $S^{?(\delta_1)}(A,B)$ and $Y^{\delta_2}(\{A,C\},B)$

We consider the network fragment from Fig. 6, consisting of node B and its mutually independent parents A and C. We assume for now that the nodes A and C remain independent as observations are being entered into the network. By conditioning on A and C, we find for the probability of b:

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

We observe that Pr(b) is a function of Pr(a) and Pr(c), and that for a fixed Pr(c), Pr(b) is a linear function of Pr(a). For Pr(a) $= 1$, the function yields Pr($b \mid a$); for Pr(a) $= 0$, it yields Pr($b \mid \bar{a}$). Moreover, the gradient of the function at a particular Pr(c) matches the sign of the situational influence of node A on node B for that Pr(c). In essence, we have two different, so-called, *manifestations* of the non-monotonic influence of A on B: either the sign of the situational influence is negative for low values of Pr(c) and positive for high values of Pr(c), as shown in Fig. 7, or vice versa, as shown in Fig. 8.

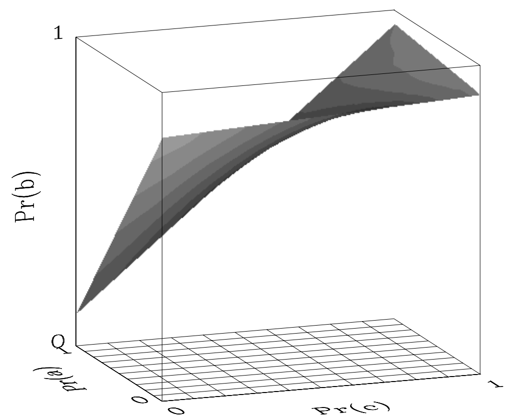

Fig. 7. An example Pr(b) as a function of Pr(a) and Pr(c), with $S^?(A, B), S^+(C, B)$ and $Y^+(\{A, C\}, B)$

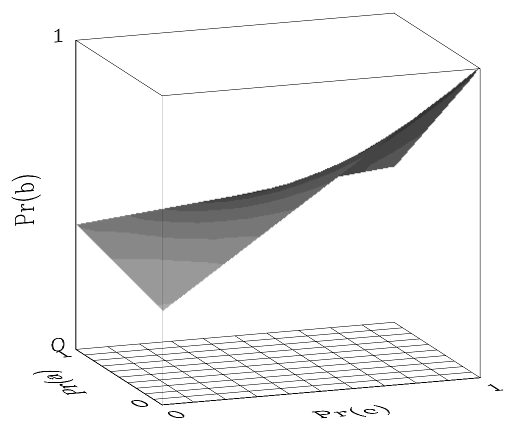

Fig. 8. An example Pr(b) as a function of Pr(a) and Pr(c), with $S^?(A, B), S^+(C, B)$ and $Y^-(\{A, C\}, B)$

As a result of observations being entered into the network, the probability of c may change. The sign of the situational influence of node A on node B may then change as well. For some changes of the probability of c, however, the sign will definitely stay the same. Whether or not it will do so depends on the manifestation of the non-monotonic influence, on the current sign, and on the direction of change of the probability of c. In the graph depicted in Fig. 7, for example, the sign of the situational influence 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 the graph depicted in Fig. 8. A method for verifying whether or not the sign of a situational influence retains its validity thus has to distinguish between the two possible manifestations of the underlying non-monotonic influence.

The sign of the additive synergy involved can now aid in distinguishing between the possible manifestations of a non-monotonic influence under study. We recall that a positive additive synergy of nodes A and C on their common child B indicates that $\Pr(b \mid ac) - \Pr(b \mid \bar{a}c) \geq \Pr(b \mid a\bar{c}) - \Pr(b \mid \bar{a}\bar{c})$. From the influence of A on B being non-monotonic, we have that the differences $\Pr(b \mid ac) - \Pr(b \mid \bar{a}c)$ and $\Pr(b \mid a\bar{c}) - \Pr(b \mid \bar{a}\bar{c})$ have opposite signs. A positive additive synergy of A and C on B now implies that the sign of $\Pr(b \mid ac) - \Pr(b \mid \bar{a}c)$ must be positive and that the sign of $\Pr(b \mid a\bar{c}) - \Pr(b \mid \bar{a}\bar{c})$ must be negative, as in Fig. 7. Analogously, a negative additive synergy corresponds to the manifestation of the non-monotonic influence shown in Fig. 8.

From the previous observations, we have that the sign of the additive synergy involved can be exploited for verifying whether or not the sign of a situational influence retains its validity during inference. Suppose that, as in Fig. 6, we have $S^{?(\delta_1)}(A, B)$ and $Y^{\delta_2}(\{A, C\}, B)$. Further suppose that new evidence causes a change in the probability of c, the direction of which is reflected in sign$[C]$. Then, we can be certain that δ_1 will remain valid if

$$\delta_1 = \text{sign}[C] \otimes \delta_2 \,.$$

Otherwise, δ_1 has to be changed into '?'. We can substantiate our statement as follows. Abstracting from previously entered evidence, we have that

$$\begin{aligned}\Pr(b \mid a) - \Pr(b \mid \bar{a}) = {} & \Pr(c) \cdot (\Pr(b \mid ac) - \Pr(b \mid a\bar{c}) - \Pr(b \mid \bar{a}c) + \Pr(b \mid \bar{a}\bar{c})) + \\ & \Pr(b \mid a\bar{c}) - \Pr(b \mid \bar{a}\bar{c})\,.\end{aligned}$$

We observe that the equation expresses the difference $\Pr(b \mid a) - \Pr(b \mid \bar{a})$ as a linear function of $\Pr(c)$. We further observe that the sign of the gradient of this function equals the sign of the additive synergy of A and C on B. Now suppose that the probability of c increases as a result of the new evidence, and that $Y^{+}(\{A, C\}, B)$. Since the gradient then is positive, a positive sign for the situational influence will remain valid. If, on the other hand, the probability of c increases and $Y^{-}(\{A, C\}, B)$, then a negative sign for the situational influence will remain valid. We conclude that upon an increase of $\Pr(c)$, δ_1 persists if $\delta_1 = + \otimes \delta_2$. Otherwise, we cannot be certain of the sign of the situational influence and δ_1 is changed to '?'. Similar observations hold for a decreasing probability of c.

In our analysis so far, we have assumed that the two parents A and C of a node 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. Node 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 by the sign-propagation algorithm building upon the composition of signs.

4.2 The Adapted Sign-Propagation Algorithm

The sign-propagation algorithm for inference with a qualitative network has to be adapted to render it applicable to situational qualitative networks. In essence, two modifications are required. First, in case of non-monotonicities, the algorithm must use the signs of the situational influences involved. Furthermore, because the sign of a situational influence of a node A on a node B is dynamic, its validity has to be verified as soon as an observation causes a change in the probability distribution of another parent of B. Due to the nature of sign propagation, it may occur that a sign is propagated along a situational influence between A and B, while the fact that the probability distribution of another parent of B changes does not become apparent until later in the propagation. It may then turn out that the sign of the situational influence should have been adapted and that incorrect signs were propagated. A solution to this problem is to verify the validity of the sign of the situational influence as soon as information to this end becomes available; if the sign requires updating, inference is restarted with the updated network. Since the sign of a situational influence can change only once, the number of restarts is limited. The adapted part of the sign-propagation algorithm is summarised in pseudo-code in Fig. 9.

Example 5. We consider the situational qualitative network from Fig. 10. The network is identical to the one shown in Fig. 4, except that it is supplemented with a situational sign for the non-monotonic influence of node A on node B. Suppose that we are again interested in the effect of observing the value *false* for node A upon the other nodes in the network. Inference starts with node A receiving the message '$-$' and updating its node sign to $0 \oplus - = -$. Node A subsequently determines the messages to be sent to its neighbours E, B and D. To node E, it sends $- \otimes - = +$. Upon receiving this message, node E updates its node sign to $0 \oplus + = +$ as before; node E does not propagate the message to B. To node B, node A sends the message $- \otimes - = +$, using the sign of the situational influence. Node B updates its node sign to $0 \oplus + = +$. It subsequently computes the message $+ \otimes + = +$ for E. Upon receiving this message, node E does not need to change its node sign. Node B does not propagate the message it has received from A to node C. To node D, node A

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 neighbour V_i of *to* in Q
  do *linksign* $\leftarrow$ sign of influence between *to* and V_i;
    message $\leftarrow$ sign[*to*] $\otimes$ *linksign*;
    if $V_i \notin$ *trail* and sign[V_i] $\neq$ sign[V_i] $\oplus$ *message*
    then Propagate-Sign(Q,*trail*,V_i,*message*).

procedure Determine-Effect-On(Q, V_i):
  $N_{V_i} \leftarrow \{V_j \rightarrow V_k \mid V_j \in \pi(V_k) \setminus \{V_i\}, V_k \in \sigma(V_i), S^{?(\delta)}(V_j, V_k), \delta \neq ?\}$;
  for all $V_j \rightarrow V_k \in N_{V_i}$
  do Verify-Update($S^{?(\delta)}(V_j, V_k)$);
  if a δ changes
  then $Q \leftarrow Q$ with adapted signs;
    return Process-Observation(Q,O,*sign*).

Fig. 9. The adapted part of the sign-propagation algorithm

sends $- \otimes + = -$. Node D updates its node sign to $0 \oplus - = -$. It subsequently determines the message $- \otimes + = -$ for node C. Upon receiving this message, C updates its node sign to $0 \oplus - = -$. The algorithm now establishes that node C is a parent of node B which has node A for its other parent, and that the influence of node A on B is non-monotonic. Because the node sign of C has changed, the validity of the sign of the situational influence of A on B needs to be verified. Since $- = - \otimes +$, the algorithm finds that the sign of the situational influence of A on B remains valid. The inference therefore continues. Node C sends the message $- \otimes - = +$ to B. Since node B need not change its node sign, the inference halts. The node signs resulting from the inference are shown in the network's nodes in Fig. 10. □

Examples 3 and 5 demonstrate that inference with a situational network can yield more informative results when compared to a regular qualitative network.

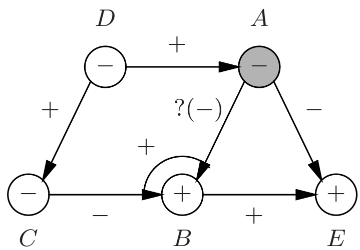

Fig. 10. A situational network and its node signs after the observation $A = false$

5 Conclusions and Further Research

Qualitative probabilistic networks model the probabilistic relationships between their variables by means of signs. If such a relationship is non-monotonic, it has associated the ambiguous sign '?', even though the influence is always unambiguous in the current state of the network. The presence of influences with ambiguous signs typically leads to ambiguous, and thus uninformative, results upon inference. In this paper we extended qualitative networks with situational influences that capture qualitative information about the current effect of non-monotonic influences. We showed that these situational influences can be used upon inference and may then effectively forestall ambiguous inference results. Because the signs of situational influences are dynamic in nature, we identified conditions under which these signs retain their validity. We studied the dynamics of the signs of situational influences in cases where the non-monotonicity involved originates from a single variable. The presented ideas and methods, however, are readily generalised to cases where the non-monotonicity is caused by more than one variable. To conclude, we adapted the existing sign-propagation algorithm to situational qualitative networks.

By introducing situational influences we have, in essence, strengthened the expressiveness of a qualitative network. Recently, other research has also focused on enhancing the formalism of qualitative networks, for example by introducing a notion of strength of influences [5]. In the future we will investigate how these different enhancements can be integrated to arrive at an even more powerful framework for qualitative probabilistic reasoning.

Acknowledgement. This research has been 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.J. Druzdzel and L.C. van der Gaag. Elicitation of probabilities for belief networks: combining qualitative and quantitative information. In Ph. Besnard and S. Hank, editors, *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pages 141–148, Morgan Kaufmann Publishers, San Francisco, 1995.
3. M.P. Wellman. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, 44:257–303, 1990.
4. M.J. Druzdzel and M. Henrion. Efficient reasoning in qualitative probabilistic networks. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 548–553, AAAI Press, Menlo Park, California, 1993.
5. S. Renooij and L.C. van der Gaag. Enhancing QPNs for trade-off resolution. In K. Laskey and H. Prade, editors, *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, pages 480–487, Morgan Kaufmann Publishers, San Francisco, 1999.