

Propagation of Multiple Observations in QPNs Revisited

Silja Renooij¹, Linda C. van der Gaag¹, and Simon Parsons²

Abstract. The sign-propagation algorithm for inference with a qualitative probabilistic network has been designed to handle a single observation at a time. Multiple observations can in essence be dealt with by entering them consecutively and combining the results of the successive propagations, or by entering them for a newly added dummy node. We demonstrate that both approaches can yield weaker results than necessary. We identify the causes underlying this unnecessary weakness and adapt the propagation algorithm so as to provide for the strongest possible results upon inference.

1 INTRODUCTION

Qualitative probabilistic networks (QPNs) were introduced in the early 1990s for probabilistic reasoning in a qualitative way [1]. A qualitative network encodes statistical variables and the probabilistic relationships between them in a directed graph. The encoded relationships in essence represent influences on probability distributions. Each such influence captures the shift in distribution for a variable that is occasioned by a shift in another variable's distribution. The direction of the occasioned shift is summarised by a qualitative sign.

Qualitative networks can play an important role in the construction of quantitative probabilistic networks for real-life applications. The construction of a quantitative network begins with the construction of its graph. Assessment of the various probabilities required, which often is very hard, is typically performed only when the network's graph is considered robust. Now, by associating signs with the relationships modelled in the graph, a qualitative network is obtained that can be used for studying the reasoning behaviour of the quantitative network prior to the assessment of its probabilities. For this purpose, it is important to derive as much information as possible from the qualitative network.

Inference with a qualitative probabilistic network is based upon the idea of combining and propagating signs [2]. The basic algorithm computes the effect of a single observation on all the variables in the network. It yields, for each variable, a sign indicating the direction of the shift in distribution that is occasioned by the new observation. In real-life applications, often the simultaneous, joint effect of multiple observations is of interest. Multiple observations can in essence be dealt with in two ways [3]. One way is to add a dummy descendant to the observed nodes, for which an appropriate observation is entered and subsequently propagated. Another way is to enter and propagate the various observations one after the other and combine the results of the successive propagations to yield the joint effect. Unfortunately, both approaches can yield weaker results than necessary.

In this paper, we address the propagation of multiple observations in a qualitative probabilistic network. We will show that the dynamics of the set of influences over which signs are propagated, can affect the results of inference. We further show that some influences are guaranteed to be dominated by others and should be disregarded during sign propagation. Building upon these properties, we adapt the basic algorithm to yield the strongest possible inference results. The paper is organised as follows. Section 2 briefly reviews qualitative networks and the basic sign-propagation algorithm. Section 3 discusses propagating multiple observations with the two approaches outlined above. In the Sections 4 and 5 we study the dynamics of influences and their dominance properties. In Section 6, we revisit the basic algorithm. The paper ends with some conclusions in Section 7.

2 QUALITATIVE PROBABILISTIC NETWORKS

A *qualitative probabilistic network* encodes statistical variables and the probabilistic relationships between them in a directed acyclic graph $G = (V(G), A(G))$. Each node $A \in V(G)$ represents a variable. For ease of exposition, we assume all variables to be binary with $a > \bar{a}$, writing a for $A = \text{true}$ and \bar{a} for $A = \text{false}$, but our results are readily generalised to non-binary variables. The set $A(G)$ of arcs captures probabilistic independence between the represented variables. We say that a chain between two nodes is *blocked* if it includes either an observed node with at least one outgoing arc, or an unobserved node with two incoming arcs and no observed descendants; a node with two incoming arcs is termed a *head-to-head node*. If all chains between two nodes are blocked, then these nodes are said to be *d-separated* and the corresponding variables are considered conditionally independent given the observations present [4].

Associated with its digraph, a qualitative probabilistic network specifies influences and synergies [1]. A *qualitative influence* between two nodes expresses how the values of one node influence the probability distribution over the values of the other node.

Definition 2.1 Let G be an acyclic digraph with $A \rightarrow B \in A(G)$. A positive qualitative influence of A on B , denoted $S^+(A, B)$, expresses that observing a higher value for A makes the higher value for B more likely, regardless of any other direct influences on B , that is, $\Pr(b \mid ax) \geq \Pr(b \mid \bar{a}x)$ for any combination of values x for the set $\pi(B) \setminus \{A\}$ of predecessors of B other than A .

The '+' in $S^+(A, B)$ is termed the *sign* of the influence. A negative qualitative influence S^- and a zero influence S^0 are defined analogously. If the influence of A on B is not monotonic or if it is unknown after inference, we say that it is *ambiguous*, denoted $S^?(A, B)$. The definition of qualitative influence is generalised straightforwardly to influences along *chains* without head-to-head nodes.

The set of influences of a qualitative network exhibits various properties. The *symmetry* property states that, if the network in-

¹ Institute of Information and Computing Sciences, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands. Email: {silja,linda}@cs.uu.nl

² Massachusetts Institute of Technology, Sloan School of Management, 3 Cambridge Center, NE20-336 Cambridge, MA 02142, USA. Email: sparsons@mit.edu

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

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

cludes 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 chain without head-to-head nodes combine into a single influence with the \otimes -operator from Table 1. The property of *composition* asserts that multiple parallel influences between two nodes combine into a single influence with the \oplus -operator.

A qualitative probabilistic network further includes *product synergies* [5], that express how the value of one node influences the probability distribution of another node given a value for a third node.

Definition 2.2 Let G be as before, with $A, B, C \in V(G)$ and $\pi(C) = \{A, B\}$. A negative product synergy of A on B (and vice versa) given the value c for node C , denoted $X^-(\{A, B\}, c)$, expresses that, given c , a higher value for A renders the higher value for B less likely, that is, $\Pr(c|ab) \cdot \Pr(c|\bar{a}\bar{b}) \leq \Pr(c|a\bar{b}) \cdot \Pr(c|\bar{a}b)$.

Positive, zero, and ambiguous synergies are defined analogously. The product synergy $X^\delta(\{A, B\}, c)$ serves, upon observing c , to induce a qualitative *intercausal* influence with sign δ between A and B .

Example 2.1 We consider the qualitative network from Figure 1, which is a simplified fragment of a real-life network in the field of oesophageal cancer. Node U represents whether or not a patient’s tumour is ulcerating; node L models whether or not the tumour is longer than 10 cm. Node W indicates whether or not the tumour has grown beyond the oesophageal wall into adjacent structures. U and L are modelled as the possible causes of tumour growth outside the oesophagus. Since the presence of either cause suffices to increase the probability of invasion of adjacent structures, both U and L exert a positive qualitative influence on W , indicated by the signs over the arcs. The network further models that either value for W induces a negative intercausal influence between U and L , indicated by the two signs over the dashed line. Given a tumour’s growth beyond the oesophageal wall, for example, the negative synergy expresses that observation of one cause explains away the other cause. \square

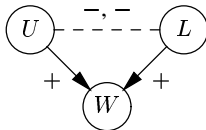


Figure 1. The qualitative *Wall invasion* network.

Inference with a qualitative network is based upon the idea of propagating and combining signs [2]. The algorithm traces the effects of observing a node’s value on the other nodes in the network by message-passing between neighbours. For each node, a *node sign* is determined, indicating the direction of shift in its probability distribution as occasioned by the new observation. Initially, all node signs equal ‘0’. For the newly observed node, an appropriate sign is entered, that is, a ‘+’ for the value *true* or a ‘-’ for *false*. Each node receiving a message updates its sign with the \oplus -operator, and then sends a message to each neighbour that is not d-separated from the

observed node and to every node on which it exerts an induced influence³. The sign of this message is the \otimes -product of the node’s (new) sign and the sign of the influence it traverses. This process is repeated throughout the network, building on the properties of symmetry, transitivity, and composition of influences. Since each node can change its sign at most twice, the algorithm is guaranteed to halt.

3 PROPAGATING MULTIPLE OBSERVATIONS

The sign-propagation algorithm for inference with a qualitative network basically serves to compute the effects of a *single* observation. The algorithm can, however, be used to handle multiple observations [3]. The first approach is to add a dummy node D to the network, with arcs $O_i \rightarrow D$ for each observed node O_i ; the sign of the influence associated with the arc $O_i \rightarrow D$ corresponds to the sign of the observation for O_i . Running the basic sign-propagation algorithm with a ‘+’ for the dummy node will now yield the joint effect of all the observations. The other approach is to enter and propagate the various observations one after the other. The joint effect on unobserved nodes then equals the sign-sum of the results of the successive propagations; observed nodes retain their sign of observation. Both approaches, unfortunately, tend to yield unnecessarily weak results.

In the dummy-node approach to handling multiple observations, the node sign of the newly added node is set to ‘+’. The node signs of the truly observed nodes, however, are not fixed and can therefore change during inference, as is illustrated by the following example.

Example 3.1 We consider the qualitative network from Figure 2. Again pertaining to the invasion of an oesophageal tumour into adjacent structures beyond the oesophagus, it describes that the longer the tumour, the more likely it is to have grown through the oesophageal wall. The length L of the tumour is strongly correlated with whether or not the tumour is circular, modelled by node C . We now address entering the observations $L = \text{true}$ and $C = \text{false}$.

Figure 2 shows the results of propagating the two observations with the dummy-node approach. First, a dummy node *Dummy* is added to the network, with arcs $L \rightarrow \text{Dummy}$ and $C \rightarrow \text{Dummy}$, and influences $S^+(L, \text{Dummy})$ and $S^-(C, \text{Dummy})$. A ‘+’ is then entered for the dummy node. *Dummy* sends a ‘+’ to node L . L thereupon sends a ‘+’ to both C and W . The dummy node also sends a ‘-’ to node C , which in turn passes the ‘-’ on to L . All nodes end up with the ambiguous node sign ‘?’ after inference. \square

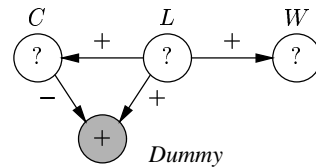


Figure 2. The effect of entering a ‘+’ for node L and a ‘-’ for node C , using the dummy-node approach.

The example illustrates a typical problem with the dummy-node approach: since multiple observations are entered as a single observation for a dummy node, the actual observations do not block chains as they would in the original network. The set of influences over which signs are propagated may thus be too large. In fact, observations can even be propagated to nodes from which they are actually

³ The literature on sign propagation is not clear on whether or not an intercausal influence that is induced by an observation is immediately used upon propagating that observation. Here we assume that induced influences are not used immediately. In Section 5 we will justify our assumption.

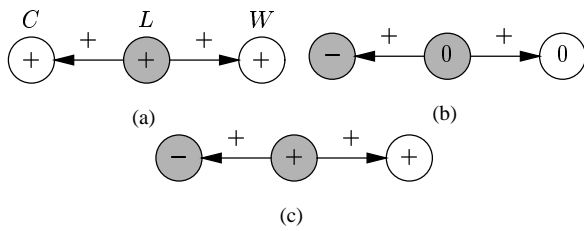


Figure 3. The separate effects of entering a '+' for L (a) and a '-' for C (b), and their joint effect (c).

d-separated. This can result in weaker signs than necessary, that is, it can result in '?'s instead of '+', '-' or '0's. Although not incorrect, these ambiguous signs are very uninformative and, moreover, tend to spread to large parts of the network.

Using the sequential-propagation approach to handling multiple observations, the order in which observations are entered can affect the net result of inference. The differences upon inference originate from the *dynamics* of the set of influences over which signs are propagated: the set shrinks as chains are blocked and expands as intercausal influences are induced. We present two illustrative examples.

Example 3.2 We consider again the network fragment from Figure 2, this time without the dummy node. We again address entering the two observations $L = true$ and $C = false$.

Figure 3 shows the results from first entering $L = true$ and then $C = false$. To propagate the first observation, node L sends a '+' to both C and W . All three nodes end up with a positive node sign. Then, a '-' is entered for node C . As node L is observed, its sign is not affected. In addition, as node L blocks the chain from C to W , no sign is passed on to node W . The joint effect of the two observations thus is positive for nodes L and W , and negative for node C . These are the strongest possible results derivable from the network.

Figure 4 shows the results from first entering $C = false$ and then $L = true$. The joint effect of these observations on the probability distribution of node W now is unknown. Note that this result is weaker than necessary. □

Example 3.3 We consider once again the network fragment from Figure 1 and address the observations $L = true$ and $W = true$.

Figure 5 shows the results from first entering $L = true$ and then $W = true$. After entering the first observation, node L propagates a '+' to node W . As U and L are independent causes of W , it does not pass on a message to U : U 's probability distribution is not affected by the observation. Then, a '+' is entered for node W . W sends a '+' to both U and L . As node L is observed, its sign is not affected by the new observation. The joint effect of the two observations shows a positive net influence on node U 's probability distribution, which is the strongest possible result that can be derived from the network.

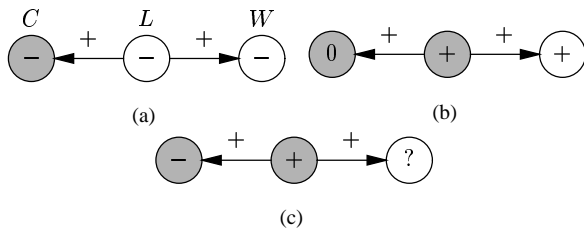


Figure 4. The separate effects of entering a '-' for C (a) and a '+' for L (b), and their joint effect (c).

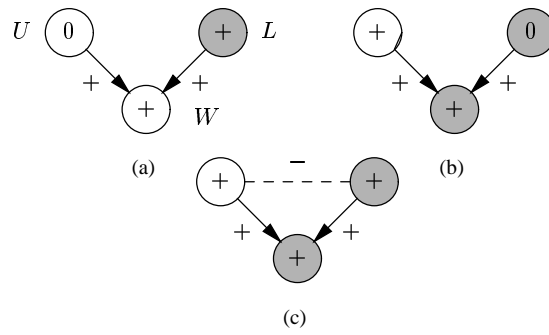


Figure 5. The separate effects of entering a '+' for L (a) and a '+' for W (b), and their joint effect (c).

Figure 6 now shows the results from first entering $W = true$ and then $L = true$. After entering the first observation, node W propagates a '+' to both its causes. The observation in addition induces a negative intercausal influence between U and L . Subsequently entering the second observation causes node L to send a '-' over the intercausal influence to node U . It further sends a '+' to W , but node W has been observed and will not change sign. The joint effect of the observations now reveals an ambiguous net effect on U 's probability distribution, which is a correct but unnecessarily weak result. □

The previous examples demonstrate that, by sequential propagation of multiple observations and adding the results, the order in which the various observations are entered can influence the results and can yield unnecessary ambiguous node signs. As mentioned before, the differences in results can be attributed to the dynamics of the set of influences over which signs are propagated.

4 EXPLOITING DYNAMICS

When a single observation is entered into a qualitative network, the basic algorithm propagates the associated sign to each node that is not d-separated from the observed node. The set of influences over which the sign is propagated is then unique. By entering multiple observations one after the other, however, this set changes dynamically. On the one hand, influences are removed as chains are blocked; on the other hand, intercausal influences are added [7]. In the previous section we have shown that these dynamics can give rise to unnecessary '?'s and can yield different results upon inference, dependent upon the order in which the observations are entered.

The order of entering observations for two nodes O_1 and O_2 can influence the sign resulting for a node A if A is d-separated from, for example, O_1 given O_2 . Then, if the observation for O_1 is entered

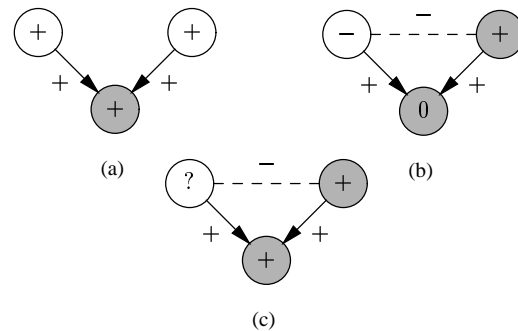


Figure 6. The separate effects of entering a '+' for W (a) and for L (b), and their joint effect (c).

first, a sign may be propagated to node A that would not have reached it if the observation for O_2 had been entered first. To ensure that the order in which multiple observations are entered is immaterial, therefore, the sign of an observation should be propagated only along chains that will not be blocked by subsequent observations. To this end, for each node O_i in the set O of simultaneously observed nodes, we determine the set of nodes that are d-separated from O_i given $O \setminus \{O_i\}$ and the set P of all previously observed nodes. We call this set the *exclusion set* $X(O_i)$ for O_i . We now have that any node in $X(O_i)$ is independent of O_i given $(O \cup P) \setminus \{O_i\}$.

Proposition 4.1 *Let G be an acyclic digraph and let \Pr be a probability distribution that respects G . Now, let P be a set of previously observed nodes and let O be the set of newly observed nodes. Then, for each $O_i \in O$, $\Pr(X(O_i) | O \cup P) = \Pr(X(O_i) | (O \cup P) \setminus \{O_i\})$.*

Now, to ensure that the order of entering multiple observations does not affect net results, the sign-propagation algorithm should restrict the propagation of an observation for $O_i \in O$ to those nodes in the digraph that are not included in $X(O_i)$. From Proposition 4.1, we have that, by doing so, no node ever receives a sign it should not have received given the other observations.

With the dummy-node approach to handling multiple observations, the set of influences over which signs are propagated does not change dynamically upon inference since only a single observation for the dummy node is entered. However, as demonstrated in the previous section, observations can then be propagated to nodes from which they are d-separated. To ensure that signs are propagated correctly, the sign-propagation algorithm should be adapted to send signs to observed nodes over influences from the dummy node only.

5 DOMINANCE OF INFLUENCES

In the previous section, we discussed the use of an exclusion set to prevent the order in which multiple observations are entered into a qualitative network from affecting the net results of inference. In this section, we demonstrate that any intercausal influences induced by the observations should be disregarded during sign propagation. To this end, we investigate the *sole* effect of these intercausal influences.

We begin by showing that the influences that are induced by a specific observation should not be used in propagating that observation.

Proposition 5.1 *Let G be the digraph of a qualitative network with $A \rightarrow C \in A(G)$ and $S^\delta(A, C)$, and without any other chains between A and C 's predecessors than through C . Let δ_C be the sign of an observation for node C and let δ_A be the sign of node A given this observation. Then, $\delta_A = \delta_C \otimes \delta$.*

Proof: We prove the proposition for $\delta = \delta_C = +$; proofs for other combinations of signs are analogous. The sign of node A given the observation for C equals the sign of the change $\Pr(a | c) - \Pr(a)$ in A 's probability distribution. We find for all combinations of values x for the set X of predecessors of C other than A , that

$$\begin{aligned} \Pr(a | cx) - \Pr(a) &= \frac{(\Pr(c | ax) - \Pr(c | x)) \cdot \Pr(a)}{\Pr(c | x)} \\ &= (\Pr(c | ax) - \Pr(c | \bar{a}x)) \cdot \left(\frac{\Pr(\bar{a}) \cdot \Pr(a)}{\Pr(c | x)} \right) \end{aligned}$$

From $S^+(A, C)$ we have that $\Pr(c | ax) - \Pr(c | \bar{a}x) \geq 0$ for all x . We conclude that $\Pr(a | c) - \Pr(a) = \sum_x (\Pr(a | cx) - \Pr(a)) \cdot \Pr(x | c) \geq 0$ and, hence, that $\delta_A = \delta_C \otimes \delta$. \square

The previous proposition only pertains to the effect of the intercausal influence on node A . Any other influences on A from C 's predecessors are handled by the propagation algorithm. We would like to note that the proposition holds also when a descendant of C is observed.

We now show that, if an observation pertains to a node from a set of multiple simultaneously observed nodes, then the intercausal influence induced by that observation should also be disregarded when propagating the *other* observations; more specifically, we show that direct influences always dominate over intercausal ones.

Dominance of direct influences over intercausal ones was already suggested by M.J. Druzdzel [3, Section 6.4.3]. Druzdzel focuses on the situation where a head-to-head node and one of its parents are observed. As we will show presently, in this situation the effect of the observation of the head-to-head node on an unobserved parent is larger than the effect of the observation for the observed parent via the induced intercausal influence. Druzdzel claims that this dominance property follows from the following proposition: "For parents A, B of C , and A, C of D , the qualitative influence of D on B solely depends on the influence of D on C and that of C on B ." In the network described, there are two chains from D to B , one consisting of $D \leftarrow A \rightarrow C \leftarrow B$ and one consisting of $D \leftarrow C \leftarrow B$. Druzdzel proves the proposition by observing that the latter chain is the only unblocked chain. Note, however, that if node D , a descendant of the head-to-head node C , is observed, then an intercausal influence is induced between nodes A and C , and the chain from A to B via C becomes unblocked as well. Unfortunately, as the proposition does not mention observed nodes nor intercausal influences, we feel that it does not correctly capture the dominance property.

We will formally show that the dominance property suggested by Druzdzel indeed holds. We say that the influence of a node B on a node C *dominates* the influence of a node A on C , if an observation for B has a larger effect on the probability distribution of C than an observation for A , that is, iff, for all observations $a_i \in \{a, \bar{a}\}$ of A and $b_i \in \{b, \bar{b}\}$ of B , we have that $|\Pr(c | b_i) - \Pr(c)| \geq |\Pr(c | a_i) - \Pr(c)|$. We now prove that direct influences dominate over intercausal ones.

Proposition 5.2 *Let G be the digraph of a qualitative network with $A \rightarrow C, B \rightarrow C \in A(G)$ and without any other chains between A and B than through C . Let δ_{direct} be the sign of change for node A given an observation for node C and let δ_{inter} be the sign of change for node A given a subsequent observation for node B . Let δ_A be the sign of node A given both observations. Then, $\delta_A = \delta_{direct}$.*

Proof: We prove the proposition for $S^+(A, C)$, the observations c for node C and b for node B , and $X^-(\{A, B\}, c)$. Note that we then have that $\delta_{direct} = +$ and $\delta_{inter} = -$. Proofs for other combinations of signs are analogous. The sign δ_{direct} of the change in A 's probability distribution equals the sign of the difference $\Pr(a | c) - \Pr(a)$. The sign δ_{inter} of the change in A 's probability distribution occasioned by the subsequent observation for node B , equals the sign of $\Pr(a | bc) - \Pr(a | c)$. The node sign δ_A of node A given both observations equals the sign of the difference $\Pr(a | bc) - \Pr(a)$. For all combinations of values x for the set X of predecessors of C other than A and B , we now have that

$$\begin{aligned} \Pr(a | bcx) - \Pr(a) &= \frac{\Pr(c | abx) \cdot \Pr(a | bx)}{\Pr(c | bx)} - \Pr(a) \\ &= \frac{(\Pr(c | abx) - \Pr(c | bx)) \cdot \Pr(a)}{\Pr(c | bx)} \\ &= (\Pr(c | abx) - \Pr(c | \bar{a}bx)) \cdot \left(\frac{\Pr(a) \cdot \Pr(\bar{a})}{\Pr(c | bx)} \right) \end{aligned}$$

From $S^+(A, C)$ we have for all values x that $\Pr(c | abx) - \Pr(c | \bar{a}bx) \geq 0$. We conclude that $\Pr(a | bc) - \Pr(a) \geq 0$ and, hence, that $\delta_A = \delta_{direct}$. \square

We note that the proposition also holds for indirect observations for node C . The dominance property tells us that during the sequential propagation of multiple simultaneous observations, intercausal influences induced by any of these observations should be disregarded.

6 PROBABILISTIC INFERENCE REVISITED

In the previous sections, we argued that observations for the nodes O_i from a set O of simultaneously observed nodes should be propagated only to the nodes that are not d-separated from O_i given all other nodes from O . In addition, we demonstrated that, upon propagating multiple simultaneous observations, we should disregard the intercausal influences induced by these observations. The basic sign-propagation algorithm can be easily adapted to incorporate these ideas. For this purpose, a node's exclusion set can be computed efficiently with the well-known *Bayes-Ball* algorithm [6]. This algorithm computes the set of nodes that are *structurally irrelevant* for a node of interest given all observed nodes. For an observed node O_i , the set of structurally irrelevant nodes given all other observations, is exactly our exclusion set $X(O_i)$ of O_i .

We now illustrate the impact of disregarding intercausal influences and using exclusion sets upon propagating multiple observations.

Example 6.1 We consider the qualitative network from Figure 7. Figure 7(a) shows the results of using the basic sign-propagation algorithm, after entering the subsequent observations $D = true$, $B = true$, $C = false$, and $G = false$, and combining the results. Note that the nodes A and E end up with the node sign '?' as a result of propagating a negative sign over the intercausal links induced by the observations for D and B . The '?' is propagated from node E to node G and onwards, before the observation for G is entered. As a result, nodes H , I and J also end up with the sign '?'. Figure 7(b) shows the results of using the adapted algorithm, which disregards

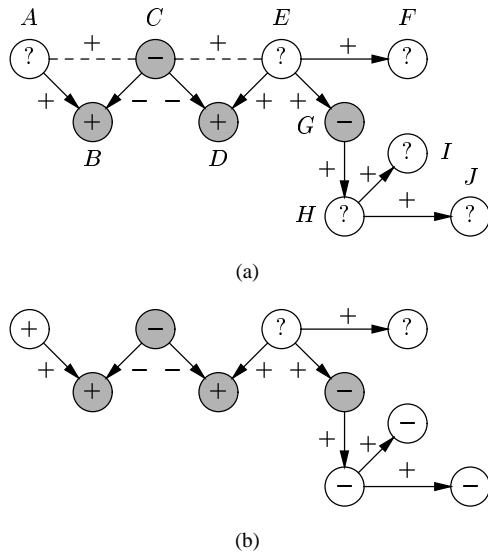


Figure 7. Propagation of multiple observations using the original sign-propagation algorithm (a) and the adapted algorithm (b).

intercausal influences and exploits the fact that node G d-separates the nodes $\{H, I, J\}$ from the other observed nodes. \square

The adapted algorithm prevents the propagation of signs to nodes that are d-separated from the observed node, given all other observed nodes from the set of multiple observations. Each node therefore receives a sign if and only if its probability distribution is truly influenced by the entire set of observations. As no node ever receives signs it should not have received, no unnecessary '?'s are generated and the algorithm returns the strongest possible signs that can be derived from the specification of the network.

7 CONCLUSIONS

The basic algorithm for probabilistic inference with a qualitative network has been designed to determine the effect of a single observation on the probability distributions of all nodes in the network. We demonstrated that handling multiple observations by applying the basic algorithm for each observation separately and combining the results into their joint effect can yield weaker results than necessary. Furthermore, the results may depend on the order in which the observations are entered. As the cause of these problems, we identified the dynamics of the set of influences over which signs are propagated upon inference. We showed that the intercausal influences that are added to this set are always dominated by direct influences and should therefore be disregarded upon inference. In addition, we showed that using exclusion sets for observed nodes can prevent the propagation of the sign of an observation to nodes that will be d-separated from the observed node given subsequently entered observations. The adapted sign-propagation algorithm yields results that are the strongest that can be derived from the qualitative network as specified.

The concept of exclusion set and the disregarding of intercausal influences are not just valuable when sequentially propagating multiple observations, but can also be exploited with the dummy-node approach. We feel however that the dummy-node approach, when compared to sequential propagation, bears the major drawback of changing the structure of a network.

ACKNOWLEDGEMENTS

This research has been (partly) supported by the Netherlands Organisation for Scientific Research (NWO).

REFERENCES

- [1] M.P. Wellman (1990). Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, vol. 44, pp. 257 – 303.
- [2] M.J. Druzdzel and M. Henrion (1993). Efficient reasoning in qualitative probabilistic networks. *Proceedings of the Eleventh National Conference on Artificial Intelligence*, Morgan Kaufmann, pp. 548 – 553.
- [3] M.J. Druzdzel (1993). *Probabilistic Reasoning in Decision Support Systems: From Computation to Common Sense*, PhD Thesis, Department of Engineering and Public Policy, CMU, Pittsburgh, PA.
- [4] J. Pearl (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Palo Alto.
- [5] M.J. Druzdzel and M. Henrion (1993). Intercausal reasoning with uninstantiated ancestor nodes. *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, pp. 317 – 325.
- [6] L.C. van der Gaag and J.-J.Ch. Meyer (1996). The dynamics of probabilistic structural relevance. *Proceedings of the Eighth Dutch Conference on Artificial Intelligence*, Utrecht University, pp. 145 – 156.
- [7] R.D. Shachter (1998). Bayes-Ball: The rational pastime (for determining irrelevance and requisite information in belief networks and influence diagrams). *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, pp. 480 – 487.