

Zooming in on Trade-offs in Qualitative Probabilistic Networks

Silja Renooij, Linda C. van der Gaag,

Dept. of Computer Science
Utrecht University
P.O. Box 80.089
3508 TB Utrecht
The Netherlands
{silja,linda}@cs.uu.nl

Shaw Green,

Dept. of Electronic Engineering
Queen Mary & Westfield College
University of London
London E1 4NS
United Kingdom
s.d.green@elec.qmw.ac.uk

Simon Parsons

Dept. of Computer Science
University of Liverpool
Chadwick Building
Liverpool L69 7ZF
United Kingdom
s.d.parsons@csc.liv.ac.uk

Abstract

Qualitative probabilistic networks have been designed for probabilistic reasoning in a qualitative way. As a consequence of their coarse level of representation detail, qualitative probabilistic networks do not provide for resolving trade-offs and typically yield ambiguous results upon inference. We present an algorithm for computing more informative results for unresolved trade-offs. The algorithm builds upon the idea of zooming in on the truly ambiguous part of a qualitative probabilistic network and identifying the information that would serve to resolve the trade-offs present.

Introduction

Qualitative probabilistic networks were introduced in the early 1990s for probabilistic reasoning with uncertainty in a qualitative way (Wellman 1990). A qualitative probabilistic network encodes variables and the probabilistic relationships between them in a directed acyclic graph. The encoded relationships basically represent influences on the variables' probability distributions. Each of these influences is summarised by a qualitative sign indicating a direction of shift in probability distribution. For probabilistic inference with qualitative networks, an elegant algorithm based upon the idea of propagating and combining signs is available (Druzdzel and Henrion 1993a).

Qualitative probabilistic networks capture the relationships between their variables at a coarse level of representation detail. These networks do therefore not provide for resolving trade-offs, that is, for establishing the net result of two or more conflicting influences on a variable's probability distribution. If trade-offs are represented in a qualitative probabilistic network, then probabilistic inference will typically yield ambiguous results. Once an ambiguity arises, it will spread throughout most of the network upon inference, even if only a very small part of the network is truly ambiguous.

The issue of dealing with trade-offs in qualitative probabilistic networks has been addressed by several researchers. S. Parsons (1995) has introduced, for example, the concept of categorical influences. A categorical

influence is either an influence that serves to increase a probability to 1 or an influence that decreases a probability to 0, regardless of any other influences, thereby resolving any trade-off in which it is involved. C.-L. Liu and M.P. Wellman (1998) have designed a method for resolving trade-offs based upon the idea of reverting to numerical probabilities whenever necessary. S. Renooij and L.C. van der Gaag (1999) have enhanced the formalism of qualitative probabilistic networks by distinguishing between strong and weak influences. Trade-off resolution during inference is then based on the idea that strong influences dominate over conflicting weak influences.

In this paper, we present a new algorithm for dealing with trade-offs in qualitative probabilistic networks. Rather than resolve trade-offs by providing for a finer level of representation detail, our algorithm identifies from a qualitative probabilistic network the information that would serve to resolve the trade-offs present. From this information, a more insightful result than ambiguity is constructed.

Our algorithm for dealing with trade-offs builds upon the idea of zooming in on the part of a qualitative probabilistic network where the actual trade-offs reside. After an observation has been entered into a network, the sign of the influence of this observation on a variable of interest is computed. If the sign is ambiguous, then there are trade-offs present in the network. In fact, a trade-off must reside along the reasoning chains between the observation and the variable of interest. Our algorithm isolates these reasoning chains to constitute the part of the network that is relevant for addressing trade-offs. From this relevant part, an informative result is constructed for the variable of interest in terms of values for the variables involved and the relative strengths of the influences among them.

The paper is organised as follows. We set out by presenting some preliminaries concerning qualitative probabilistic networks. We then introduce the basic idea of our algorithm for zooming in on trade-offs informally, by means of an example. The algorithm is thereupon discussed in further detail. The paper ends with some concluding observations.

Preliminaries

A *qualitative probabilistic network* encodes statistical variables and the probabilistic relationships between them in a directed acyclic graph. Each node in the digraph represents a variable. Each arc can be looked upon as expressing a causal influence from the node at the tail of the arc on the node at the arc's head. More formally, the digraph's set 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. 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 entered observations (Pearl 1988).

A qualitative probabilistic network associates with its digraph qualitative influences and qualitative synergies (Wellman 1990). A *qualitative influence* between two nodes expresses how the values of one node influence the probabilities of the values of the other node. A positive qualitative influence of node A on its successor B expresses that observing higher values for A makes higher values for B more likely, regardless of any other direct influences on B ; the influence is denoted $S^+(A, B)$, where '+' is the influence's *sign*. A negative qualitative influence, denoted S^- , and a zero qualitative influence, denoted S^0 , are defined analogously. If the influence of node A on node B is not monotonic or unknown, we say that it is *ambiguous*, denoted $S^?(A, B)$.

The set of influences of a qualitative probabilistic network exhibits various properties (Wellman 1990). 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 property of *transitivity* asserts that qualitative influences along a chain that specifies at most one incoming arc for each node, combine into a single influence with the \otimes -operator from Table 1. The property of *composition* asserts that multiple influences between two nodes along parallel chains combine into a single influence with the \oplus -operator.

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

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

In addition to influences, a qualitative probabilistic network includes *synergies* that express how the value of one node influences the probabilities of the values of another node in view of a value for a third node (Druzdzel and Henrion 1993b). A negative product synergy of node A on node B (and vice versa) given the value c for their common successor C , denoted $X^-(\{A, B\}, c)$, expresses that, given c , higher values for A render higher values for B less likely. A product synergy induces

a qualitative influence between the predecessors of a node upon observation; the induced influence is coined an *intercausal influence*. Positive, zero, and ambiguous product synergies are defined analogously.

Example 1 We consider the small qualitative probabilistic network shown in Figure 1. The network rep-




Figure 1: The qualitative *antibiotics* network.

resents a fragment of fictitious and incomplete medical knowledge, pertaining to the effects of administering antibiotics on a patient. Node A represents whether or not a patient takes antibiotics. Node T models whether or not a patient has typhoid fever and node D represents presence or absence of diarrhoea. Node F describes whether or not the composition of a patient's bacterial flora has changed.

Typhoid fever and a change in bacterial flora are modelled as the possible causes of diarrhoea. As the presence of either of them will increase the probability of a patient having diarrhoea, the influences of both T and F on D are positive. Antibiotics can cure typhoid fever by killing the bacteria that cause the infection; the influence of A on T , therefore, is negative. Antibiotics can also change the composition of a patient's bacterial flora, thereby increasing the risk of diarrhoea; the influence of A on F is positive. Upon observing diarrhoea in a patient, the presence of typhoid fever in itself is a sufficient explanation, reducing the probability that a change in bacterial flora is also a contributing cause; a similar observation holds for a change in composition of bacterial flora. Given diarrhoea, therefore, a negative intercausal influence is induced between T and F .

The qualitative *antibiotics* network models two conflicting influences on the probability distribution of node D and therefore captures a trade-off. For a patient who is known to take antibiotics, the trade-off cannot be resolved and the result with regard to this patient having diarrhoea is ambiguous. \square

For inference with a qualitative network, an elegant algorithm is available from M.J. Druzdzel and M. Henrion (1993a). The basic idea of the algorithm is to trace the effect of observing a node's value on the other nodes in a network by message-passing between neighbouring nodes. For each node, a *node sign* is determined, indicating the direction of change in the node's probability distribution occasioned by the new observation given all previously observed node values. Initially, all node signs equal '0'. For the newly observed node, an appropriate sign is entered, that is, either a '+' for the observed value *true* or a '-' for the value *false*. Each

node receiving a message updates its sign and subsequently 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 intercausal influence. The sign of this message is the \otimes -product of the node's (new) sign and the sign of the influence it traverses. This process is repeated throughout the network, building on the properties of symmetry, transitivity, and composition of influences. Each node is visited at most twice, since a node can change sign at most twice, and the process is therefore guaranteed to halt.

Outline of the Algorithm

If a qualitative probabilistic network models trade-offs, it will typically yield ambiguous results upon inference with the sign-propagation algorithm. From Table 1, we have that whenever two conflicting influences on a node are combined with the \oplus -operator, an ambiguous sign will result. Once an ambiguous sign is introduced, it will spread throughout most of the network and an ambiguous sign is likely to result for the node of interest. By zooming in on the part of the network where the actual trade-offs reside and identifying the information that would serve to resolve them, a more insightful result can be constructed. We illustrate the basic idea of our algorithm for this purpose.

As our running example, we consider the qualitative probabilistic network from Figure 2. Now, suppose that the value *true* has been observed for the node *H* and that we are interested in its influence on the probability distribution of node *A*. Tracing the influence of the observation on every node's distribution by means of the basic sign-propagation algorithm, results in the node signs as shown in Figure 3. These signs reveal that at least one trade-off must reside along the reasoning chains between the observed node *H* and the node of interest *A*. These chains together constitute the part of the network that is relevant for addressing the trade-offs that have given rise to ambiguous results; this part is termed the *relevant network*. For the example, the




Figure 2: The example qualitative network.




Figure 3: The result of propagating '+' for node *H*.

relevant network is shown in Figure 4 below the dashed line. Our algorithm isolates this relevant network for further investigation. To this end, it deletes from the network all nodes and arcs that are connected to, but no part of the reasoning chains from *H* to *A*.

A relevant network for addressing trade-offs typically includes many nodes with ambiguous node signs. Often, however, only a small number of these nodes are actually involved in the trade-offs that have given rise to ambiguous results. Figure 4, for example, reveals that, while the nodes *A*, *B*, and *C* have ambiguous node signs, the influences between them are not conflicting. In fact, any unambiguous node sign $sign[C]$ for node *C* would result in the unambiguous node sign $sign[C] \otimes ((+ \otimes -) \oplus -) = sign[C] \otimes -$ for node *A*. For addressing the trade-offs involved, therefore, the part of the relevant network between node *C* and node *A* can be disregarded. Node *C* is termed the *pivot node* for the node of interest *A*. In general, the pivot node is a node with an ambiguous sign for which *every possible* unambiguous sign would uniquely




Figure 4: The relevant network, below the dashed line.

determine an unambiguous sign for the node of interest; in addition, the pivot node does not reside on an unblocked chain from another node having this property to the node of interest, that is, the pivot node is the node with this property “closest” to the observed node. Our algorithm now computes from the relevant network the pivot node for the network’s node of interest.

From the definition of pivot node, we have that there must be two or more reasoning chains from the observed node to the pivot node; the net influence along these chains must be conflicting. Our algorithm identifies the information that would serve to resolve the ambiguity at the pivot node. For this purpose, the algorithm selects a minimal set of nodes, each with two or more incoming arcs, for which unambiguous node signs would uniquely determine the signs of the separate influences on the pivot node. These nodes with each other constitute the *resolution frontier* for the pivot node. In terms of signs for these nodes, the algorithm now constructs a sign for the pivot node by comparing the relative strengths of its various conflicting reasoning chains.




Figure 5: The construction of a sign for node C .

In the example network, two influences are exerted on the pivot node C : the influence from node F via node D on C and the influence from E on C . Note that unambiguous signs for the nodes F and E would render both influences unambiguous. These nodes with each other now are taken to constitute the resolution frontier for node C . For the sign δ of the influence of node F via node D on C and for the sign δ' of the influence of E on C , we find that

$$\begin{aligned} \delta &= \text{sign}[F] \otimes \delta_1 \otimes \delta_3 & \delta' &= \text{sign}[E] \otimes \delta_4 \\ &= \text{sign}[F] \otimes + & &= \text{sign}[E] \otimes - \end{aligned}$$

where δ_i , $i = 1, 3, 4$, are as in Figure 5. For the node sign $\text{sign}[C]$ of the pivot node, the algorithm now constructs the following result:

$$\text{if } |\delta| \geq |\delta'|, \text{ then } \text{sign}[C] = \delta, \text{ else } \text{sign}[C] = \delta'$$

where $|\delta|$ denotes the strength of the sign δ . So, if the two influences on node C have opposite signs, then their relative strengths will determine the sign for node C . The sign of the node of interest A then follows directly from the sign of C .

Splitting up and Constructing Signs

In this section we further detail some of the issues involved in our algorithm for zooming in on trade-offs. In doing so, we assume that a qualitative network does not include any ambiguous influences, that is, ambiguous node signs upon inference result from unresolved

trade-offs. We also assume that a single observation is entered into the network and that sign-propagation results in an ambiguous sign for the node of interest. We focus attention on identifying the pivot node from the relevant part of a qualitative network and on constructing an informative result for the network’s node of interest; further details are provided in a forthcoming technical paper.

Splitting up the Network

Our algorithm identifies from a qualitative network the relevant part for addressing the trade-offs that have resulted in an ambiguous sign for the node of interest. From the relevant network, the pivot node is identified.

The relevant network is constructed by reducing the original network’s digraph. First, the *computationally relevant* part of the network is identified. In a *quantitative* probabilistic network, a node is said to be computationally relevant to a node of interest, if its (conditional) probability distribution is required for computing the posterior probability distribution for this node of interest given all previously observed nodes. For computing the set of computationally relevant nodes, the efficient *Bayes-Ball* algorithm is available from R.D. Shachter (1998). From the computationally relevant network, all nodes are identified that do not reside on any reasoning chain from the newly observed node to the node of interest; these nodes are removed to yield the relevant network. An efficient algorithm is available from Y. Lin and M.J. Druzdzel (1997) to identify these so-called *nuisance nodes*.

From the relevant network, the pivot node is identified. We recall that the pivot node is a node with an ambiguous sign for which any unambiguous sign would uniquely determine an unambiguous sign for the node of interest. From this property, we have that the pivot node is either the node of interest or an *articulation node* in the relevant network. An articulation node is a node that upon removal, along with its incident arcs, makes the digraph fall apart into various components; articulation nodes are found by depth-first search (Cormen, Leiserson, and Rivest 1990). Our algorithm now sets out by computing all articulation nodes in the relevant network. As any reasoning chain in the relevant network from the observed node to the node of interest visits all articulation nodes, we have that there exists a total ordering on these nodes. Numbering them from 1, closest to the observed node, to m , closest to the node of interest, the pivot node basically is the articulation node with the lowest number for which an unambiguous sign would uniquely determine an unambiguous sign for the node of interest. To identify the pivot node, our algorithm starts with the articulation node numbered m and investigates whether an unambiguous sign for this node would result in an unambiguous sign for the node of interest upon sign propagation. If the sign of the node of interest is ambiguous, then the node of interest itself is the pivot node. Note that, in the qualitative *antibiotics* network from Figure 1, the node of interest

is the pivot node. Otherwise, the algorithm proceeds by investigating the articulation node numbered $m - 1$, and so on.

Constructing Results

From its definition, we have that the pivot node for a qualitative network's node of interest receives two or more conflicting net influences and, hence, captures a trade-off. Our algorithm now focuses on this trade-off and identifies the information that would serve to resolve it. For this purpose, our algorithm computes the so-called *candidate resolvers* for the pivot node. A candidate resolver is a node with an ambiguous node sign that has two or more incoming arcs and resides on a chain from the observed node to the pivot node. From among these candidate resolvers, a minimal set of nodes is constructed for which unambiguous node signs would uniquely determine the signs of the separate influences on the pivot node. This set of nodes constitutes the so-called *resolution frontier*. The resolution frontier is computed to be the set of candidate resolvers that do not reside on a chain from another candidate resolver to the pivot node. In terms of signs for the nodes from the resolution frontier, the algorithm now constructs an informative result for the pivot node by comparing the relative strengths of the various influences upon it.

Let F be the resolution frontier for the pivot node P . For each resolver $R_i \in F$, let $sign[R_i]$ be its node sign. Let s_j^i , $j \geq 1$, denote the signs of the different reasoning chains from R_i to the pivot node. For each combination of node signs $sign[R_i]$, $R_i \in F$, the sign of the pivot node is computed to be

$$\text{if } \left| \oplus_{(sign[R_i] \otimes s_j^i)=+} (sign[R_i] \otimes s_j^i) \right| \geq \left| \oplus_{(sign[R_i] \otimes s_j^i)=-} (sign[R_i] \otimes s_j^i) \right|$$

then $sign[P] = +$, else $sign[P] = -$

where $|\delta|$ once again denotes the strength of the sign δ . The process of thus constructing informative results can be repeated recursively for the pivot node's resolvers.

Conclusions

We have presented a new algorithm for dealing with trade-offs in qualitative probabilistic networks. Rather than resolve trade-offs by providing for a finer level of representation detail, our algorithm identifies from a qualitative network the information that would serve to resolve the trade-offs present. For this purpose, the algorithm zooms in on the ambiguous part of the network and identifies the pivot node for the node of interest. For the pivot node, a more informative result than ambiguity is constructed in terms of values for the node's resolvers and the relative strengths of the influences upon it. This process of constructing informative results can be repeated recursively for the pivot node's resolvers.

We believe that qualitative probabilistic networks can play an important role in the construction of Bayesian

belief networks for real-life application domains. The construction of a Bayesian belief network typically sets out with the construction of the network's digraph. As the assessment of the various probabilities required is a far harder task, it is performed only when the network's digraph is considered robust. Now, by assessing signs for the influences modelled in the digraph, a qualitative network is obtained that can be exploited for studying the projected belief network's reasoning behaviour prior to the assessment of probabilities. For this purpose, algorithms are required that serve to derive as much information as possible from a qualitative network. We look upon our algorithm as a first step in this direction.

Acknowledgment

This work was partially supported by the EPSRC under grant GR/L84117 and a PhD studentship.

References

- Cormen, T.H., Leiserson, C.E., and Rivest R.L. 1990. *Introduction to Algorithms*. Cambridge, Massachusetts: MIT Press.
- Druzdzel, M.J., and Henrion, M. 1993a. Efficient reasoning in qualitative probabilistic networks. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, 548 – 553.
- Druzdzel, M.J., and Henrion, M. 1993b. Intercausal reasoning with uninstantiated ancestor nodes. In *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence*, 548 – 553.
- Lin, Y., and Druzdzel, M.J. 1997. Computational advantages of relevance reasoning in Bayesian belief networks. In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, 342 – 350.
- Liu, C.-L., and Wellman, M.P. 1998. Incremental tradeoff resolution in qualitative probabilistic networks. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, 338 – 345.
- Parsons, S. 1995. Refining reasoning in qualitative probabilistic networks. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, 427 – 434.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Palo Alto: Morgan Kaufmann Publishers.
- Renooij, S., and Van der Gaag, L.C. 1999. Enhancing QPNs for trade-off resolution. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, 559 – 566.
- Shachter, R.D. 1998. Bayes-Ball: the rational pastime (for determining irrelevance and requisite information in belief networks and influence diagrams). In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, 480-487.
- Wellman, M.P. 1990. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence* 44: 257 – 303.