Context-specific Sign-propagation in Qualitative Probabilistic Networks

Silja Renooij[1], **Simon Parsons**[2,*], and **Linda C. van der Gaag**[1]
[1]Institute of Information and Computing Sciences, Utrecht University
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
{silja, linda}@cs.uu.nl
[2]Department of Computer Science, University of Liverpool
Chadwick Building, Liverpool L69 7ZF, United Kingdom
s.d.parsons@csc.liv.ac.uk

Abstract

Qualitative probabilistic networks represent probabilistic influences between variables. Due to the level of representation detail provided, knowledge about influences that hold only in specific contexts cannot be expressed. The results computed from a qualitative network, as a consequence, can be quite weak and uninformative. We extend the basic formalism of qualitative probabilistic networks by providing for the inclusion of context-specific information about influences and show that exploiting this information upon inference has the ability to forestall unnecessarily weak results.

1 Introduction

Qualitative probabilistic networks are qualitative abstractions of probabilistic networks [Wellman, 1990], introduced for probabilistic reasoning in a qualitative way. A qualitative probabilistic network encodes statistical variables and the probabilistic relationships between them in a directed acyclic graph. Each node A in this digraph represents a variable. An arc $A \rightarrow B$ expresses a probabilistic influence of the variable A on the probability distribution of the variable B; the influence is summarised by a qualitative sign indicating the direction of shift in B's distribution. For probabilistic inference with a qualitative network, an efficient algorithm, based upon the idea of propagating and combining signs, is available [Druzdzel & Henrion, 1993].

Qualitative probabilistic networks can play an important role in the construction of probabilistic networks for real-life application domains. While constructing the digraph of a probabilistic network is doable, the assessment of all probabilities required is a much harder task and is only performed when the network's digraph is considered robust. By eliciting signs from domain experts, the obtained qualitative probabilistic network can be used to study and validate the reasoning behaviour of the network prior to probability assessment; the signs can further be used as constraints on the probabilities to be assessed [Druzdzel & Van der Gaag, 1995]. To be able to thus exploit a qualitative probabilistic network, it should capture as much qualitative information from the application domain as possible. In this paper, we propose an extension to the basic formalism of qualitative networks to enhance its expressive power for this purpose.

Probabilistic networks provide, by means of their digraph, for a qualitative representation of the conditional independences that are embedded in a joint probability distribution. The digraph in essence captures independences between nodes, that is, it models independences that hold for all values of the associated variables. The independences that hold only for specific values are not represented in the digraph but are captured instead by the conditional probabilities associated with the nodes in the network. Knowledge of these latter independences allows further decomposition of conditional probabilities and can be exploited to speed up inference. For this purpose, a notion of *context-specific independence* was introduced for probabilistic networks to explicitly capture independences that hold only for specific values of variables [Boutilier *et al.*, 1996; Zhang & Poole, 1999].

A qualitative probabilistic network equally captures independences between variables by means of its digraph. Since its qualitative influences pertain to variables as well, independences that hold only for specific values of the variables involved cannot be represented. In fact, qualitative influences implicitly *hide* such context-specific independences: if the influence of a variable A on a variable B is positive in one context, that is, for one combination of values for some other variables, and zero in all other contexts – indicating independence – then the influence is captured by a positive sign. Also, positive and negative influences may be hidden: if a variable A has a positive influence on a variable B in some context and a negative influence in another context, then the influence of A on B is modelled as being ambiguous.

As context-specific independences basically are qualitative by nature, we feel that they can and should be captured explicitly in a qualitative probabilistic network. For this purpose, we introduce a notion of *context-specific sign*. We extend the basic formalism of qualitative networks by providing for the inclusion of context-specific information about influences and show that exploiting this information upon inference can prevent unnecessarily weak results. The paper is organised as follows. In Section 2, we provide some preliminaries concerning qualitative probabilistic networks. We present two examples of the type of information that can be hidden in

*This work was partly funded by the EPSRC under grant GR/L84117

qualitative influences, in Section 3. We present our extended formalism and associated algorithm for exploiting context-specific information in Section 4. In Section 5, we discuss the context-specific information that is hidden in the qualitative abstractions of two real-life probabilistic networks. In Section 6, we briefly show that context-specific information can also be incorporated in qualitative probabilistic networks that include a qualitative notion of strength of influences. The paper ends with some concluding observations in Section 7.

2 Qualitative probabilistic networks

A *qualitative probabilistic network* models statistical variables as nodes in its digraph; from now on, we use the terms variable and node interchangeably. We assume, without loss of generality, that all variables are binary, using a and $\bar{a}$ to indicate the values *true* and *false* for variable A, respectively. A qualitative network further associates with its digraph a set of *qualitative influences*, describing probabilistic relationships between the variables [Wellman, 1990]. A qualitative influence associated with an arc $A \rightarrow B$ expresses how the values of node A influence the probabilities of the values of node B. A *positive qualitative influence*, for example, of A on B, denoted $S^{+}(A, B)$, expresses that observing higher values for node A makes higher values for node B more likely, regardless of any other 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 X of parents of B other than A. The '$+$' in $S^{+}(A, B)$ is termed the influence's *sign*. A negative qualitative influence S^{-}, and a zero qualitative influence S^{0}, are defined analogously. If the influence of node A on node B is non-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 *symmetry* property states that, if $S^{\delta}(A, B)$, then also $S^{\delta}(B, A)$, $\delta \in \{+, -, 0, ?\}$. The *transitivity* property asserts that a sequence of qualitative influences along a chain that specifies at most one incoming arc per node, combine into a single influence with the $\otimes$-operator from Table 1. The *composition* property asserts that multiple influences between two nodes along parallel chains combine into a single influence with the $\oplus$-operator.

$\otimes$	$+$	$-$	0	$?$
$+$	$+$	$-$	0	$?$
$-$	$-$	$+$	0	$?$
0	0	0	0	0
$?$	$?$	$?$	0	$?$

$\oplus$	$+$	$-$	0	$?$
$+$	$+$	$?$	$+$	$?$
$-$	$?$	$-$	$-$	$?$
0	$+$	$-$	0	$?$
$?$	$?$	$?$	$?$	$?$

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

A qualitative network further captures *qualitative synergies* between three or more nodes; for details we refer to [Druzdzel & Henrion, 1993; Wellman, 1990].

For inference with a qualitative network, an efficient algorithm is available [Druzdzel & Henrion, 1993]. The basic idea of the algorithm is to trace the effect of observing a node's value on the other nodes in the 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 node sign and subsequently sends a message to each neighbour whose sign needs updating. 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, once from '0' to '$+$' or '$-$', and then only to '?', the process visits each node at most twice and is therefore guaranteed to halt.

3 Context-independent signs

Context-specific information cannot be represented explicitly in a qualitative probabilistic network, but is hidden in the network's qualitative influences. If, for example, the influence of a node A on a node B is positive for one combination of values for the set X of B's parents other than A, and zero for all other combinations of values for X, then the influence of A on B is positive by definition. The zero influences are hidden due to the fact that the inequality in the definition of qualitative influence is not strict. We present an example illustrating such hidden zeroes.

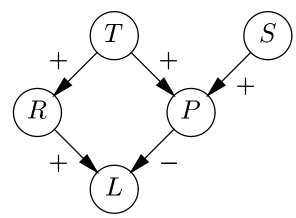

Figure 1: The qualitative *surgery* network.

Example 1 The qualitative network from Figure 1 represents a highly simplified fragment of knowledge in oncology; it pertains to the effects and complications to be expected from treatment of oesophageal cancer. Node L models the life expectancy of a patient after therapy; the value l indicates that the patient will survive for at least one year. Node T models the therapy instilled; we consider surgery, modelled by t, and no treatment, modelled by $\bar{t}$, as the only alternatives. The effect to be attained from surgery is a radical resection of the oesophageal tumour, modelled by node R. After surgery a life-threatening pulmonary complication, modelled by node P, may result; the occurrence of this complication is heavily influenced by whether or not the patient is a smoker, modelled by node S.

We consider the conditional probabilities from a quantified network representing the same knowledge. We would like to note that these probabilities serve illustrative purposes only; although not entirely unrealistic, they have not been specified by domain experts. The probability of attaining a radical resection upon surgery is $\Pr(r \mid t) = 0.45$; as without surgery there can be no radical resection, we have $\Pr(r \mid \bar{t}) = 0$.

From these probabilities we have that node T indeed exerts a positive qualitative influence on node R. The probabilities of a pulmonary complication occurring and of a patient's life expectancy after therapy are, respectively,

$\Pr(p)$	s	$\bar{s}$
t	0.75	0.00
$\bar{t}$	0.00	0.00

$\Pr(l)$	p	$\bar{p}$
r	0.15	0.95
$\bar{r}$	0.03	0.50

From the left table, we verify that both T and S exert a positive qualitative influence on node P. The fact that the influence of T on P is actually zero in the context of the value $\bar{s}$ for node S, is not apparent from the influence's sign. Note that this zero influence does not arise from the probabilities being zero, but rather from their having the same value. From the right table we verify that node R exerts a positive influence on node L; the qualitative influence of P on L is negative. □

The previous example shows that the level of representation detail of a qualitative network can result in information hiding. As a consequence, unnecessarily weak answers may result upon inference. For example, from the probabilities involved we know that performing surgery on a non-smoker has a positive influence on life expectancy. Due to the conflicting reasoning chains from T to L in the qualitative network, however, entering the observation t for node T will result in a '?' for node L, indicating that the influence is unknown.

We recall from the definition of qualitative influence that the sign of an influence of a node A on a node B is independent of the values for the set X of parents of B other than A. A '?' for the influence of A on B may therefore hide the information that node A has a positive influence on node B for some combination of values of X and a negative influence for another combination. If so, the ambiguous influence is *non-monotonic* in nature and can in fact be looked upon as specifying different signs for different contexts. We present an example to illustrate this observation.

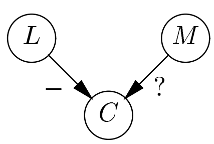

Figure 2: The qualitative *cervical metastases* network.

Example 2 The qualitative network from Figure 2 represents another fragment of knowledge in oncology; it pertains to the metastasis of oesophageal cancer. Node L represents the location of the primary tumour that is known to be present in a patient's oesophagus; the value l models that the tumour resides in the lower two-third of the oesophagus and the value $\bar{l}$ expresses that the tumour is in the oesophagus' upper one-third. An oesophageal tumour upon growth typically gives rise to lymphatic metastases, the extent of which are captured by node M. The value $\bar{m}$ of M indicates that just the local and regional lymph nodes are affected; m denotes that distant lymph nodes are affected. Which lymph nodes are local or regional and which are distant depends on the location of the tumour in the oesophagus. The lymph nodes in the neck, or cervix, for example, are regional for a tumour in the upper one-third of the oesophagus and distant otherwise. Node C represents the presence or absence of metastases in the cervical lymph nodes.

We consider the conditional probabilities from a quantified network representing the same knowledge; once again, these probabilities serve illustrative purposes only. The probabilities of the presence of cervical metastases in a patient are

$\Pr(c)$	l	$\bar{l}$
m	0.35	0.95
$\bar{m}$	0.00	1.00

From these probabilities we have that node L indeed has a negative influence on node C. The influence of node M on C, however, is non-monotonic:

$$\Pr(c \mid ml) > \Pr(c \mid \bar{m}l), \text{ yet } \Pr(c \mid m\bar{l}\,) < \Pr(c \mid \bar{m}\bar{l}\,)$$

The non-monotonic influence hides a '+' for the value l of node L and a '−' for the context $\bar{l}$. □

From the two examples above, we observe that context-specific information about influences that is present in the conditional probabilities of a quantified network cannot be represented explicitly in a qualitative probabilistic network: upon abstracting the quantified network to the qualitative network, the information is effectively hidden.

4 Context-specificity and its exploitation

The level of representation detail of a qualitative probabilistic network enforces influences to be independent of specific contexts. In this section we present an extension to the basic formalism of qualitative networks that allows for associating context-specific signs with qualitative influences. In Section 4.1, the extended formalism is introduced; in Section 4.2, we show, by means of the example networks from the previous section, that exploiting context-specific information can prevent unnecessarily weak results upon inference.

4.1 Context-specific signs

Before introducing context-specific signs, we define a notion of context for qualitative networks. Let X be a set of nodes, called the *context nodes*. A *context* c_X for X is a combination of values for a subset $Y \subseteq X$ of the set of context nodes. When $Y = \varnothing$, we say that the context is *empty*, denoted ϵ; when $Y = X$, we say that the context is *maximal*. The set of all possible contexts for X is called the *context set* for X and is denoted $\mathcal{C}_X$. To compare different contexts for the same set of context nodes X, we use an ordering on contexts: for any two combinations of values c_X and c'_X for $Y \subseteq X$ and $Y' \subseteq X$, respectively, we say that $\mathrm{c}_X > \mathrm{c}'_X$ iff $Y \supset Y'$ and c_X and c'_X specify the same combination of values for Y'.

A *context-specific sign* now basically is a sign that may vary from context to context. It is defined as a function $\delta : \mathcal{C}_X \to \{+, -, 0, ?\}$ from a context set $\mathcal{C}_X$ to the set of basic signs, such that for any two contexts c_X and c'_X with $\mathrm{c}_X > \mathrm{c}'_X$ we have that, if $\delta(\mathrm{c}'_X) = \delta_i$ for $\delta_i \in \{+, -, 0\}$, then $\delta(\mathrm{c}_X) \in \{\delta_i, 0\}$. For abbreviation, we will write $\delta(X)$ to denote the context-specific sign δ that is defined on the context set $\mathcal{C}_X$. Note that the basic signs from regular qualitative networks can be looked upon as context-specific signs that are defined by a constant function.

In our extended formalism of qualitative networks, we assign context-specific signs to influences. We say that a node A exerts a *qualitative influence of sign* $\delta(X)$ on a node B, denoted $S^{\delta(X)}(A, B)$, where X is the set of parents of B other than A, iff for each context c_X for X we have that

- $\delta(\mathrm{c}_X) = +$ iff $\Pr(b \mid a\mathrm{c}_X y) \geq \Pr(b \mid \bar{a}\mathrm{c}_X y)$ for any combination of values $\mathrm{c}_X y$ for X;
- $\delta(\mathrm{c}_X) = -$ iff $\Pr(b \mid a\mathrm{c}_X y) \leq \Pr(b \mid \bar{a}\mathrm{c}_X y)$ for any such combination of values $\mathrm{c}_X y$;
- $\delta(\mathrm{c}_X) = 0$ iff $\Pr(b \mid a\mathrm{c}_X y) = \Pr(b \mid \bar{a}\mathrm{c}_X y)$ for any such combination of values $\mathrm{c}_X y$;
- $\delta(\mathrm{c}_X) = ?$ otherwise.

Note that we take the set of parents of node B other than A for the set of context nodes; the definition is readily extended to apply to arbitrary sets of context nodes, however. Context-specific qualitative synergies can be defined analogously.

A context-specific sign $\delta(X)$ in essence has to specify a basic sign from $\{+, -, 0, ?\}$ for each possible combination of values in the context set $\mathcal{C}_X$. From the definition of $\delta(X)$, however, we have that it is not necessary to explicitly indicate a basic sign for every such context. For example, consider an influence of a node A on a node B with the set of context nodes $X = \{D, E\}$. Suppose that the sign $\delta(X)$ of the influence is defined as

$$\begin{array}{cccc} & \delta(\epsilon) = ?, & & \\ \delta(d) = +, & \delta(\bar{d}) = -, & \delta(e) = \ ?, & \delta(\bar{e}) = +, \\ \delta(de) = +, & \delta(d\bar{e}) = +, & \delta(\bar{d}e) = -, & \delta(\bar{d}\bar{e}) = 0 \end{array}$$

The function $\delta(X)$ is uniquely described by the signs of the smaller contexts whenever the larger contexts are assigned the same sign. The function is therefore fully specified by

$$\delta(\epsilon) = ?,\ \delta(d) = +,\ \delta(\bar{d}) = -,\ \delta(\bar{e}) = +,\ \delta(\bar{d}\bar{e}) = 0$$

The sign-propagation algorithm for probabilistic inference with a qualitative network, as discussed in Section 2, is easily extended to handle context-specific signs. The extended algorithm propagates and combines *basic signs* only. Before a sign is propagated over an influence, it is investigated whether or not the influence's sign is context-specific. If so, the currently valid context is determined from the available observations and the basic sign specified for this context is propagated; if none of the context nodes have been observed, then the sign specified for the empty context is propagated.

4.2 Exploiting context-specific signs

In Section 3 we presented two examples showing that the influences of a qualitative probabilistic network can hide context-specific information. Revealing this hidden information and exploiting it upon inference can be worthwhile. The information that an influence is zero for a certain context can be used, for example, to improve the runtime of the sign-propagation algorithm because propagation of a sign can be stopped as soon as a zero influence is encountered. More importantly, however, exploiting the information can prevent conflicting influences arising during inference. We illustrate this observation by means of an example.

Example 3 We reconsider the qualitative *surgery* network from Figure 1. Suppose that a non-smoker is undergoing surgery. In the context of the observation $\bar{s}$ for node S, propagating the observation t for node T with the basic sign-propagation algorithm results in the sign '?' for node L: there is not enough information present in the network to compute a non-ambiguous sign from the two conflicting reasoning chains from T to L.

We now extend the qualitative *surgery* network by assigning the context-specific sign $\delta(S)$, defined by

$$\delta(s) = +,\ \delta(\bar{s}) = 0,\ \ \delta(\epsilon) = +$$

to the influence of node T on node P, that is, we explicitly include the information that non-smoking patients are not at risk for pulmonary complications after surgery. The thus extended network is shown in Figure 3(a). We now reconsider our non-smoking patient undergoing surgery. Propagating the observation t for node T with the extended sign-propagation algorithm in the context of $\bar{s}$ results in the sign '$(+ \otimes +) \oplus (0 \otimes -)$' = '+' for node L: we find that surgery is likely to increase life expectancy for the patient. $\Box$

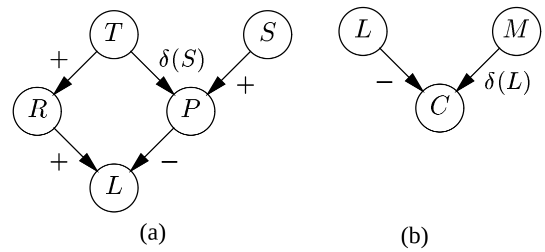

Figure 3: A hidden zero revealed, (a), and a non-monotonicity captured, (b), by a context-specific sign.

In Section 3 we not only discussed hidden zero influences, but also argued that positive and negative influences can be hidden in non-monotonic influences. As the initial '?'s of these influences tend to spread to major parts of a network upon inference, it is worthwhile to resolve the non-monotonicities involved whenever possible. Our extended formalism of qualitative networks provides for effectively capturing information about non-monotonicities, as is demonstrated by the following example.

Example 4 We reconsider the qualitative *cervical metastases* network from Figure 2. We recall that the influence of node M on node C is non-monotonic since

$$\Pr(c \mid ml) > \Pr(c \mid \bar{m}l) \text{ and } \Pr(c \mid m\bar{l}) < \Pr(c \mid \bar{m}\bar{l})$$

In the context l, therefore, the influence is positive, while it is negative in the context $\bar{l}$. In the extended network, shown in Figure 3(b), this information is captured explicitly by assigning the sign $\delta(L)$, defined by

$$\delta(l) = +,\ \delta(\bar{l}) = -,\ \delta(\epsilon) = ?$$

to the influence of node M on node C. $\Box$

5 Context-specificity in real-life networks

To get an impression of the context-specific information that is hidden in real-life qualitative probabilistic networks, we

	# *influences with sign* δ:				
	$+$	$-$	0	?	*total*
ALARM	17	9	0	20	46
oesophagus	32	12	0	15	59

Table 2: The numbers of influences with '$+$', '$-$', '0' and '?' signs for the qualitative ALARM and oesophagus networks.

computed qualitative abstractions of the well-known ALARM-network and of the network for oesophageal cancer. The ALARM-network consists of 37, mostly non-binary, nodes and 46 arcs; the number of direct qualitative influences in the abstracted network – using the basic definition of qualitative influence – therefore equals 46. The oesophagus network consists of 42, also mostly non-binary, nodes and 59 arcs. Table 2 summarises for the two abstracted networks the numbers of direct influences with the four different basic signs.

The numbers reported in Table 2 pertain to the basic signs of the qualitative influences associated with the arcs in the networks' digraphs. Each such influence, and hence each associated basic sign, covers a number of maximal contexts. For a qualitative influence associated with the arc $A \to B$, the number of maximal contexts equals 1 (the empty context) if node B has no other parents than A; otherwise, the number of maximal contexts equals the number of possible combinations of values for the set of parents of B other than A. For every maximal context, we computed the proper (context-specific) sign from the original quantified network. Table 3 summarises the number of context-specific signs covered by the different basic signs in the two abstracted networks. From the table we have, for example, that the 17 qualitative influences with sign '$+$' from the ALARM network together cover 59 different maximal contexts. For 38 of these contexts, the influences are indeed positive, but for 21 of them the influences are actually zero.

		# c_X *with sign* δ':				
ALARM		$+$	$-$	0	?	*total*
	$+$	38	–	21	–	59
δ:	$-$	–	40	11	–	51
	0	–	–	–	–	0
	?	34	24	12	28	108
	total	72	64	44	28	218

		# c_X *with sign* δ':				
oesophagus		$+$	$-$	0	?	*total*
	$+$	74	–	8	–	82
δ:	$-$	–	36	8	–	44
	0	–	–	–	–	0
	?	6	3	2	38	49
	total	80	39	18	38	175

Table 3: The numbers of contexts c_X covered by the '$+$', '$-$', '0' and '?' signs and their associated context-specific signs, for the qualitative ALARM and oesophagus networks.

For the qualitative ALARM-network, we find that 35% of the influences are positive, 17% are negative, and 48% are ambiguous; the network does not include any explicitly specified zero influences. For the extended network, using context-specific signs, we find that 32% of the qualitative influences are positive, 31% are negative, 20% are zero, and 17% remain ambiguous. For the qualitative oesophagus network, we find that 54% of the influences are positive, 21% are negative, and 25% are ambiguous; the network does not include any explicit zero influences. For the extended network, using context-specific signs, we find that 46% of the qualitative influences are positive, 22% are negative, 10% are zero, and 22% remain ambiguous.

We observe that for both the ALARM and the oesophagus network, the use of context-specific signs serves to reveal a considerable number of zero influences and to substantially decrease the number of ambiguous influences. Similar observations were made for qualitative abstractions of two other real-life probabilistic networks, pertaining to Wilson's disease and to ventricular septal defect, respectively. We conclude that by providing for the inclusion of context-specific information about influences, we have effectively extended the expressive power of qualitative probabilistic networks.

6 Extension to enhanced networks

The formalism of *enhanced qualitative probabilistic networks* [Renooij & Van der Gaag, 1999], introduces a qualitative notion of strength of influences into qualitative networks. We briefly argue that the notions from the previous sections can also be used to provide for the inclusion and exploitation of context-specific information about such strengths.

In an enhanced qualitative network, a distinction is made between strong and weak influences by partitioning the set of all influences into two disjoint subsets in such a way that any influence from the one subset is stronger than any influence from the other subset; to this end a *cut-off value* α is used. For example, a *strongly positive qualitative influence* of a node A on a node B, denoted $S^{++}(A, B)$, expresses that

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

for any combination of values x for the set X of parents of B other than A; a *weakly positive qualitative influence* of A on B, denoted $S^{+}(A, B)$, expresses that

$$0 \leq \Pr(b \mid ax) - \Pr(b \mid \bar{a}x) \leq \alpha$$

for any such combination of values x. The sign '$+_?$' is used to indicate a positive influence whose relative strength is ambiguous. Strongly negative qualitative influences S^{--}, and weakly negative qualitative influences S^{-}, are defined analogously; a negative influence whose relative strength is ambiguous is denoted $S^{-_?}$. Zero qualitative influences and ambiguous qualitative influences are defined as in regular qualitative probabilistic networks. Renooij & Van der Gaag (1999) also provide extended definitions for the $\oplus$- and $\otimes$-operators to apply to the double signs. These definitions cannot be reviewed without detailing the enhanced formalism, which is beyond the scope of the present paper; it suffices to say that the result of combining signs is basically as one would intuitively expect.

Our notion of context-specific sign can be easily incorporated into enhanced qualitative probabilistic networks. A context-specific sign now is defined as a function $\delta : \mathcal{C}_X \to \{++, +_?, +, -, -_?, --, 0, ?\}$ from a context set $\mathcal{C}_X$ to the extended set of basic signs, such that for any two contexts c_X and c'_X with $c_X > c'_X$ we have that, if the sign is strongly positive for c'_X, then it must be strongly positive for c_X, if the sign is weakly positive for c'_X, then it must be either weakly positive or zero for c_X, and if it is ambiguously positive for c'_X, then it may be (strongly, weakly or ambiguously) positive, or zero for c_X. Similar restrictions hold for negative signs. Context-specific signs are once again assigned to influences, as before.

For distinguishing between strong and weak qualitative influences in an enhanced network, a cut-off value α has to be chosen in such a way that, basically, for *all* strong influences of a node A on a node B we have that $|\Pr(b \mid ax) - \Pr(b \mid \bar{a}x)| \geq \alpha$ for *all* contexts x, and for *all* weak influences we have that $|\Pr(b \mid ax) - \Pr(b \mid \bar{a}x)| \leq \alpha$ for *all* such contexts. If, for a specific cut-off value α, there exists an influence of node A on node B for which there are contexts x and x' with $|\Pr(b \mid ax) - \Pr(b \mid \bar{a}x)| > \alpha$ and $|\Pr(b \mid ax') - \Pr(b \mid \bar{a}x')| < \alpha$, then signs of ambiguous strength would be introduced into the enhanced network, which would seriously hamper the usefulness of exploiting a notion of strength. A different cut-off value had better be chosen, by shifting α towards 0 or 1. Unfortunately, α may then very well end up being 0 or 1. The use of context-specific information about qualitative strengths can now forestall the necessity of shifting the cut-off value, as is illustrated in the following example.

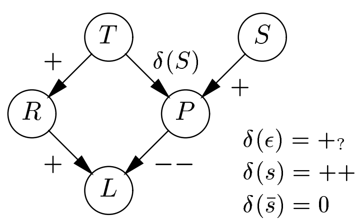

Figure 4: Context-specific sign in an enhanced network.

Example 5 We reconsider the *surgery* network and its associated probabilities from Example 1. Upon abstracting the network to an enhanced qualitative network, we distinguish between strong and weak influences by choosing a cut-off value of, for example, $\alpha = 0.46$. We then have that a pulmonary complication after surgery strongly influences life expectancy, that is, $S^{--}(P, L)$. For this cut-off value, however, the influence of node T on node P is neither strongly positive nor weakly positive; the value $\alpha = 0.46$ therefore does not serve to partition the set of influences in two distinct subsets. To ensure that all influences in the network are either strong or weak, the cut-off value should be either 0 or 1.

For the influence of node T on node P, we observe that, for $\alpha = 0.46$, the influence is strongly positive for the value s of node S and zero for the context $\bar{s}$. By assigning the context-specific sign $\delta(S)$ defined by

$$\delta(s) = ++,\ \delta(\bar{s}) = 0,\ \delta(\epsilon) = +_?$$

to the influence of node T on node P, we explicitly specify the otherwise hidden strong and zero influences. The thus extended network is shown in Figure 4. We recall from Example 3 that for non-smokers the effect of surgery on life expectancy is positive. For smokers, however, the effect could not be unambiguously determined. From the extended network in Figure 4, we now find the effect of surgery on life expectancy for smokers to be negative: upon propagating the observation t for node T in the context of the information s for node S, the sign '$(+ \otimes +) \oplus (++ \otimes --)$' = '$-$' results for node L. $\square$

7 Conclusions

We extended the formalism of qualitative probabilistic networks with a notion of context-specificity. By doing so, we enhanced the expressive power of qualitative networks. While in a regular qualitative network, zero influences as well as positive and negative influences can be hidden, in a network extended with context-specific signs this information is made explicit. Qualitative abstractions of some real-life probabilistic networks have shown that networks indeed can incorporate considerable context-specific information. We further showed that incorporating the context-specific signs into enhanced qualitative probabilistic networks that include a qualitative notion of strength renders even more expressive power. The fact that zeroes and double signs can be specified context-specifically allows them to be specified more often, in general. We showed that exploiting context-specific information about influences and about qualitative strengths can prevent unnecessary ambiguous node signs arising during inference, thereby effectively forestalling unnecessarily weak results.

References

[Boutilier *et al.*, 1996] C. Boutilier, N. Friedman, M. Goldszmidt, and D. Koller. Context-specific independence in Bayesian networks. *Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence*, 1996, pp. 115 – 123.

[Druzdzel & Henrion, 1993] M.J. Druzdzel and M. Henrion. Efficient reasoning in qualitative probabilistic networks. *Proceedings of the Eleventh National Conference on Artificial Intelligence*, 1993, pp. 548 – 553.

[Druzdzel & Van der Gaag, 1995] M.J. Druzdzel and L.C. van der Gaag. Elicitation of probabilities for belief networks: combining qualitative and quantitative information. *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, 1995, pp. 141 – 148.

[Renooij & Van der Gaag, 1999] S. Renooij and L.C. van der Gaag. Enhancing QPNs for trade-off resolution. *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, 1999, pp. 559 – 566.

[Wellman, 1990] M.P. Wellman. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, vol. 44, 1990, pp. 257 – 303.

[Zhang & Poole, 1999] N.L. Zhang and D. Poole. On the role of context-specific independence in probabilistic inference. *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, 1999, pp. 1288 – 1293.