A Case Study in the Qualitative Verification and Debugging of Numerical Uncertainty

Simon Parsons[1,2] and Alessandro Saffiotti[3]

[1]Imperial Cancer Research Fund,
London, United Kingdom.

[2]Queen Mary and Westfield College,
London, United Kingdom.

[3]IRIDIA, Université Libre de Bruxelles,
Bruxelles, Belgium.

April 18, 1996

Abstract

Quantitative methods for reasoning under uncertainty have become well established, and many alternative formalisms have been suggested. In recent years there has been a growing interest in qualitative methods as helpful in situtations in which the use of precise numerical methods is not appropriate. In this paper we demonstrate another use for qualitative models. The qualitative analysis of a quantitative model of uncertainty will reveal the qualitative behaviour of that model when new evidence is obtained. This qualitative behaviour may be studied to identify those situations in which the model does not behave as expected, and which quantitative values must be altered to correct this behaviour. The demonstration is set within the context of the diagnosis of faults in an electricity network, and reports the results of the verification of a model representing a small fragment of a real application. The model was built using Pulcinella, a tool based on Shenoy and Shafer's valuation systems.

1 INTRODUCTION

The handling of uncertainty within artificial intelligence systems has long been recognised as a topic worthy of investigation. Over two decades of research has resulted in the availability of a large number of formalisms intended for modelling different aspects of uncertainty. This work has dealt largely with complex quantitative models such as probability theory [1], possibility theory [2, 3], and evidence theory [4]. More recently, however, there has been considerable interest in the qualitative representation of reasoning under uncertainty in networks, including qualitative probabilistic networks [5, 6, 7, 8] as well as qualitative possibilistic and evidential networks [9, 10].

At the same time there has been an increase in interest in the problem of finding a unified approach to handling uncertainty, and the view that the different formal models provide exclusive approaches has been challenged by a number of authors, including Szolovits and Pauker [11], Fox [12], Saffiotti [13], Krause and Clark [14] and Neapolitan [15]. Instead the view that these formalisms are alternatives from which the most appropriate may be taken in any given situation is advanced. Such a position has been supported by the development of a generalised approach to propagating uncertainty measures in networks [16, 17], from which a general purpose tool, known as Pulcinella, for propagating uncertainty by local computation has been developed [18]. This tool has been used [19, 20] to investigate the

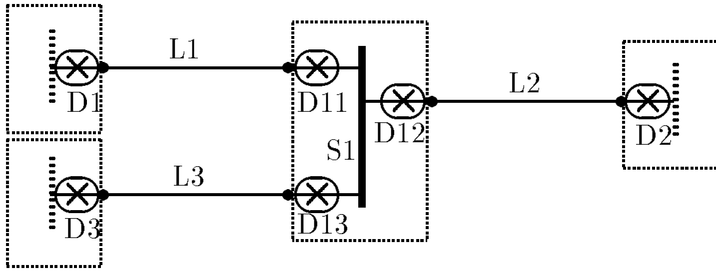

Figure 1: The relevant fragment of the distribution network.

modelling of various reasoning problems using different uncertainty handling formalisms, thus providing a comparison of the different ways in which the formalisms behave, and adding to our knowledge of their relative usefulness for solving particular problems.

In this paper we extend two of the lines of work mentioned above, by using qualitative methods [9] to advance the work of Saffiotti and Umkehrer [19], and by analysing the qualitative behaviour of the different formalisms when they are used to represent a particular problem. This information can then be used, in conjunction with expert opinion as to how the models of the problem should behave, to detect anomalies in those models, and, if any are found, to correct them. The structure of the rest of the paper is as follows. Section 2 describes the problem which we are using in this case study, along with the models in various formalisms suggested in previous work [19, 20]. Section 3 then draws upon previous results to suggest how the models may be qualitatively analysed. Section 4 uses the results of the qualitative analysis to verify the correctness of the models, and to detect some apparent anomalies. Finally Section 5 demonstrates how anomalies can be corrected, again making use of the qualitative analysis. Thus we suggest a procedure for verification and debugging:

1. Establish the qualitative behaviour of an uncertain model, whatever formalism is being used.
2. Verify that this behaviour is as required and detect any anomalies that occur.
3. Alter the model, if necessary, to ensure that the behaviour is correct.

and illustrate its use on a real problem.

2 PROBLEM DESCRIPTION

This section describes a problem in the domain of electricity distribution, reported in previous work [19, 20]. We present the three solutions originally proposed [19], and illustrate our method of qualitative debugging by analysing these solutions.

2.1 Domain Knowledge

The problem under study is to adequately model the uncertainty present in fault diagnosis in electricity networks. For the sake of clarity, both the structure of the problem and its quantitative knowledge have been greatly simplified. We consider here the fragment of an electricity network given in Figure 1.

This fragment comprises four substations, linked by three lines $L1$, $L2$ and $L3$ (in the real network, the three outer substations would be connected to other lines, substations, and so on). The substation in the middle includes $S1$, a big conductive bar, known as a busbar, used for connecting lines together. The Dis and $D1i$s are circuit breakers, that is devices which watch the part of the network on their "hot" side, marked by a dot in the picture, for overloads. If an overload is detected on a line, a circuit breaker isolates this line and transmits an alarm to the control room. The alarm may be either an instantaneous alarm or a delayed alarm. An instantaneous alarm is generated when a big overload is detected, and is normally caused by a fault in the line that the device is monitoring. A delayed alarm is generated by a small overload normally caused in an adjacent line to that being monitored. However, a delayed alarm may also be caused by an overload in the line being monitored if the overload occurs at a considerable distance from the circuit breaker. Overloads in the busbar $S1$ are detected in just the same way as faults in lines, that is an alarm from the outer circuit breakers $D1$–$D3$ may indicate a fault in $S1$. Note that a single fault can generate alarms in several circuit breakers.

Talking to domain experts revealed that the behaviour of the circuit breakers exhibits uncertainty, which has been roughly quantified by the experts:

1. alarms are not very reliable: in roughly 10% of the cases, they do not correspond to the real situation so that alarms are generated without faults or faults occur without alarms;

2. if an instantaneous alarm is generated (correctly) on a circuit breaker, the fault is in the line that the breaker is on;

3. if delayed alarm is generated (correctly) on a circuit breaker, the fault is in the line that the breaker is on in roughly 30% of cases, and in an adjacent line in roughly 70%.

In the fragment we are modelling, the "adjacent line" is the busbar $S1$ for the outer circuit breakers (i.e., $D1$, $D2$, $D3$); while it is some line outside our fragment for the inner circuit breakers ($D11$, $D12$, $D13$). Accordingly, we know what the qualitative behaviour of the modelled fragment should be (no matter what formalism is used to model the uncertainty):

1. an instantaneous alarm in an outer circuit breaker should increase our belief in the occurrence of a fault in the line that the breaker is on;

2. a delayed alarm in an outer circuit breaker should increase our belief in the occurrence of a fault in either the line the breaker is on, or in the busbar;

3. an alarm (of any kind) in an inner circuit breaker should only increase our belief in the occurrence of a fault in the line the breaker is on (in the real network, a delayed alarm would also indicate a fault in another part of the network).

2.2 Modelling the Problem

We model our problem using Shenoy and Shafer's valuation system formalism [16, 17] which is a general approach to network-based local computation, in which many uncertainty handling formalisms can be embedded. The tool we have used for our experiment, Pulcinella[1], is an implementation of valuation systems that exibits the same generality. Thanks to this generality, we have been able to use quantitative models of uncertainty as different as probability theory, possibility theory, and the Dempster-Shafer theory of belief functions. As according to the valuation system formalism, we model our problem through a set of variables, and a set of valuations linking sets of related variables. Figure 2 shows a graphical representation of the model, where ovals stand for variables, and rectangles for valuations.

A valuation over a set of variables expresses information about the values taken by the variables in that set, in a form that depends upon the uncertainty formalism; typically, this information is either a relation between those variables, or prior information about some variable. In the experiment reported here, we have modelled uncertain relations using three alternative formalisms: probability theory, possibility theory [2, 3], and the Dempster-Shafer theory of belief functions [4].[2] The Dis and $D1i$s are

[1] Pulcinella [18] was developed at IRIDIA, Université Libre de Bruxelles, and is freely available for non-commercial use. More details are available on the Pulcinella Web Page, `http://iridia.ulb.ac.be/pulcinella/`

[2] A quick introduction to the use of each of these formalisms in the context of the experiment described here can be found elsewhere [20].

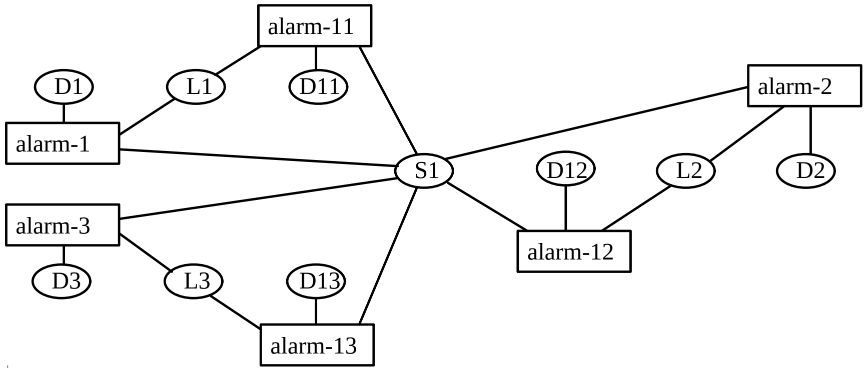

Figure 2: The valuation system for the distribution network.

variables representing circuit breaker states, with possible values ok (no alarm), del (delayed alarm), and inst (instantaneous alarm); Lis and $S1$ represent line and busbar states, with possible values ok and fault; and the "$alarm$-is" relate generation of alarms by breakers with states of neighbouring lines. Once new information about the state of a breaker is received, this can be propagated through the "$alarm$-is" relations to produce updated estimates of the states of the various elements of the network (see [16] for more details on the propagation mechanism).

In order to build the "$alarm$-i" valuations so that they behave as described above, we first split them into two groups: those referring to outer circuit breakers ($alarm$-1, $alarm$-2, $alarm$-3), and those referring to inner circuit breakers ($alarm$-11, $alarm$-12, $alarm$-13). We then enter the corresponding values. These values are meant to encode the three informal items of expert knowledge presented in Section 2.1 above. Tables 1–3 show the values for the two classes of valuations in each formalism.[3] In the probability case (Table 1), we use joint probability distributions. (For the user's convenience, these distributions are given un-normalized to Pulcinella, and then normalized by the program.) The incomplete information given by the experts had to be integrated by "reasonable" assumptions: for example, in the first column of the inner breaker case, the information "if no alarm is received, then $L1$ is ok with 90% probability" has been split equally between the possibilities "$L1 =$ ok and $S1 =$ fault" and "$L1 =$ ok and $S1 =$ ok" (as the inner circuit breakers do not give any information about possible faults in the busbar). And in the case of outer breakers, the information about the 90% reliability of alarms, and that about the relative probability of faults in $L1$ and $S1$ given a delayed alarm, have been combined into the distributions shown in the second column. It appears that filling a probability table given the available expert knowledge may require a good deal of artistic creation!

Inner breakers ($D1i$)			$P(\cdot)$		Outer breakers (Di)		
ok	del	inst	$L1$	$S1$	ok	del	inst
0.45	0.05	0.05	ok	ok	0.89	0.1	0.05
0.45	0.05	0.05	ok	fault	0.05	0.6	0.05
0.05	0.89	0.89	fault	ok	0.05	0.2	0.89
0.05	0.01	0.01	fault	fault	0.001	0.1	0.01

Table 1: The joint probability distributions for the $alarm$-i valuations

The values are much less surprising for the possibilistic case, that is, when we use possibility distributions as valuations. This is mainly due to the looser contraints imposed by possibility theory. For instance, the values in the last column of the inner case (Table 2) can be read as follows: "When

[3]The reader may not agree that this is the "best" way to model the given problem. This, of course, is not the point here, our aim being to illustrate the use of qualitative techniques to analyse and debug a possibly unsound model. To this respect, the interest of the proposed model is that it could be regarded—and indeed it has been—by a knowledge engineer as a plausible model for the given problem.

there is an instantaneous alarm, the possibility that $L1$ be ok is extremely weak, whatever the state of $S1$; on the other hand, it is completely possible that $L1$ is faulty, both with $S1$ ok or faulty".

Inner breakers ($D1i$)			$\Pi(\cdot)$		Outer breakers (Di)		
ok	del	inst	$L1$	$S1$	ok	del	inst
1	0.1	0.1	ok	ok	1	0.1	0.1
1	0.1	0.1	ok	fault	0.1	0.7	0.3
0.1	1	1	fault	ok	0.1	0.3	1
0.1	1	1	fault	fault	0.1	1	1

Table 2: The joint possibility distributions for the *alarm-i* valuations

Finally, in the belief function case (Table 3), we represent the bits of our information as multiple basic mass assignments, which are then combined by Pulcinella into one mass assignment using Dempster's rule[4]. The three mass assignments used for the outer circuit breakers are meant to encode the following beliefs, which are a way to interpret the expert's statements reported in Section 2.1:

- a 0.9 belief that when there is no alarm the lines are ok, and vice-versa; this is captured by assigning a 0.9 mass to the set of all possible configurations except those where we receive an alarm but the lines are both ok, and those where there is a fault but we do not receive an alarm (since we are ignorant about how to distribute the remaining 0.1 we follow the usual procedure for belief functions and allocate it to the full set of hypotheses expressing the fact that everything is possible to some degree).
- a 0.7 belief that if there is an instantaneous alarm, $L1$ is faulty; this is expressed by a 0.7 mass on the set of all configurations, except those where we receive an instantaneous alarm but $L1$ is not faulty (along with 0.3 to the full set, again due to ignorance). And
- a 0.3 belief that when there is a delayed alarm, then either $L1$ or $S1$ is faulty (i.e., we only exclude the case $D1 = \mathsf{del}$ & $L1 = \mathsf{ok}$ & $S1 = \mathsf{ok}$).

Things are similar for the inner breakers, except that there is no way to discriminate between $S1$ being ok or faulty—hence we omit the third mass assignment.

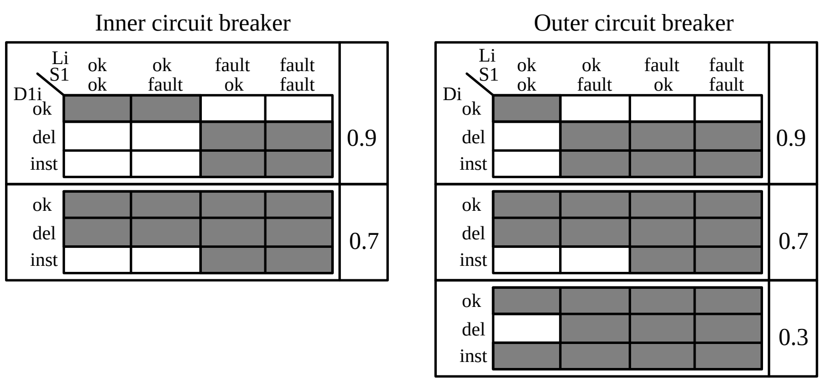

Table 3: The joint belief distributions for the *alarm-i* relations

We also have prior values for state of the breakers before any report is received (Table 4). In the case of possibility theory and of belief functions, these values indicate total ignorance about the possible state of the breaker (everything is possible, but nothing is positively believed); that is, we are agnostic

[4]The combined assignments are fairly intricate, and are not shown (but see [19]).

about whether certain alarms (or no alarms) are a priori more likely than others[5]. This contrasts with the probability case, where the axiom $p(x) + p(\neg x) = 1$ forces us to commit to $\neg x$ whatever credibility is not committed to x; here, we use a rough estimate of the priors given by the experts.

	ok	inst	del
$p(\cdot)$	0.998	0.001	0.001
$\Pi(\cdot)$	1	1	1
$bel(\cdot)$	0	0	0

Table 4: Prior values

A comparison of these three methods for handling the uncertainty in the problem is given by elsewhere [19, 20]. In the rest of this paper we will investigate whether these models of uncertainty correctly encode the behaviour described in Section 2.1.

3 QUALITATIVE ANALYSIS OF THE PROBLEM

When we find new evidence about the state of the circuit breakers we update our prior values for fault hypotheses to take account of the evidence. This takes the form of finding new values for, say, the probability of a line fault given that we have an increase in the probability of a delayed alarm (because we have evidence that one has occured). Thus we determine the change in probability of a line fault given the change in probability of a delayed alarm. When using probability, or any other theory of uncertainty, we are interested in the new value obtained after updating. When verifying the behaviour of the model, however, we are interested in checking that it corresponds to that described by the domain expert whose knowledge is captured in the model. Thus, since the expert's knowledge is often expressed in the form "if we observe A, then B is more likely", we may be more interested in knowing the way in which the values change than in the values themselves. Techniques of sensitivity analysis are sometimes used to provide this sort of verification for parametric models. Unfortunately, these techniques tend to be fairly complex, and they have only been extensively studied for the case of probabilistic models. In this section, we establish a qualitative method for sensitivity analysis that can be applied to probability, possibility and belief-function models. That is we provide a method by which it is possible to determine in a qualitative way how the values of hypotheses change when we have new evidence.

The basic method underlying the analysis has been expounded elsewhere ([9], and for binary variables [10]), and here we merely apply the results obtained there. However, for those unfamiliar with this work, we provide the following sketch. Given the equations that relate two uncertainty values val_1 and val_2 we can establish an expression, in terms of numerical uncertainty values, for the derivative $dval_1/dval_2$ that relates the two quantities. This expression allows us to determine the qualitative value of the derivative (indicated by the use of square brackets), that is whether the derivative is positive, negative, or zero. A positive derivative, written as $[dval_1/dval_2] = [+]$, means that val_1 increases when val_2 increases and decreases when val_2 decreases. Similarly, a negative derivative, $[dval_1/dval_2] = [-]$, means that val_1 decreases when val_2 increases, and increases when val_2 decreases. Finally, a zero derivative, $[dval_1/dval_2] = [0]$ means that val_1 does not change when val_2 increases or decreases. Thus when we have new evidence that changes the value of val_2 the derivative tells us the change in val_1, and for a full analysis we need to establish $[dval(h)/dval(e)]$ for every interesting hypothesis h and piece of evidence e.

[5]The last row is the result of using a vacuous basic mass assignment as prior, in other words one where all the mass is given to the full set.

3.1 The Representation for the Qualitative Analysis

The valuation system representation introduced above is that in which the original problem was formulated and solved [19, 20], and it is possible to perform the qualitative analysis directly on this representation since it contains the relevant data. However, since our previous work on qualitative analysis of uncertainty handling formalisms has been discussed using a causal network representation, it is helpful to reformulate the problem in such a representation. We use a network representation, similar to that of Pearl [1] with the exception that the numerical value of the dependencies are encoded in possibility [21] and evidence [22] theories as well as probability theory. In these networks, two nodes representing particular variables are joined by an arc if and only if the value of the variable represented by the node at one end of the arc is directly influenced by the value of the variable represented by the node at the other end of the arc. The direction of the arc is normally taken to represent the direction of the causal influence between the variables. Thus the problem information of Section 2.1 may be represented by the network of Figure 3.

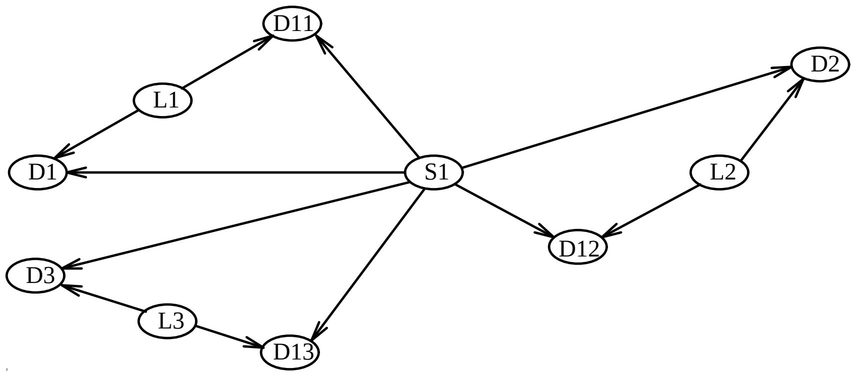

Figure 3: The causal network representation of the electricity distribution problem.

In the qualitative analysis we will look at changes in value of $S1$ and Li given changes in value of Di or $D1i$. If we wanted to assess the impact of several alarms we could sum the impact of the individual alarms using qualitative arithmetic [23].

3.2 The Probability Case

Figure 3 gives the causal network relating the causes of the different types of alarm, viz. faults in the lines and busbars, and their effects, viz. the alarms generated by the breakers. We can analyse this network to relate changes in the probabilities of alarms to the probabilities of faults. In particular, we can write the partial derivatives of the probability of a line fault by the probability of the different alarms [9], for instance:

$$\frac{\partial p(Li = \mathsf{fault})}{\partial p(D = \mathsf{inst})} = p(Li = \mathsf{fault} \mid D = \mathsf{inst}) \tag{1}$$

where D stands for both Di and $D1i$. Analagous results may be obtained for other alarm conditions. Now, this derivative tells us how $p(Li = \mathsf{fault})$ varies with $p(D = \mathsf{inst})$ ignoring the effects of changes in the probability of other values of D. We can also determine the total derivatives which tell us, for instance, how $p(Li = \mathsf{fault})$ varies when $p(D = \mathsf{inst})$ changes taking into account the changes in $p(D = \mathsf{del})$ and $p(D = \mathsf{ok})$ [9, 10]:

$$\left[\frac{dp(Li = \mathsf{fault})}{dp(D = \mathsf{inst})}\right] = \quad (2)$$
$$[p(Li = \mathsf{fault} \mid D = \mathsf{inst}) - p(Li = \mathsf{fault} \mid D = \mathsf{del})]$$
$$\oplus [p(Li = \mathsf{fault} \mid D = \mathsf{inst}) - p(Li = \mathsf{fault} \mid D = \mathsf{ok})]$$

Note that this result assumes that changes in $p(D = \mathsf{inst})$ are distributed among $p(D = \mathsf{del})$ and $p(D = \mathsf{ok})$ as a result of the fact that all three values sum to one. Again, analagous results may be obtained for other alarm states. Now, since:

$$p(Li = \mathsf{ok}) + p(Li = \mathsf{fault}) = 1 \quad (3)$$

for any line, when we know how $p(Li = \mathsf{fault})$ changes we can tell how $p(Li = \mathsf{ok})$ changes. These results are true for both inner and outer circuit breakers, and furthermore are true for any numerical values we put into the model. We can repeat the calculation for the probability of a busbar failure obtaining analagous results.

The overall change in probability of line and busbar faults due to a change in the probability of an alarm from a single circuit breaker can be calculated either from the partial derivatives, or the total derivatives. For example, the change in probability of a fault in $L1$ given a change in the state of $D1$ may be calculated as:

$$\begin{aligned}[\Delta p(L1 = \mathsf{fault})] &= \left[\frac{\partial p(L1 = \mathsf{fault})}{\partial p(D1 = \mathsf{inst})}\right] \otimes [\Delta p(D1 = \mathsf{inst})] \\ &\oplus \left[\frac{\partial p(L1 = \mathsf{fault})}{\partial p(D1 = \mathsf{del})}\right] \otimes [\Delta p(D1 = \mathsf{del})] \\ &\oplus \left[\frac{\partial p(L1 = \mathsf{fault})}{\partial p(D1 = \mathsf{ok})}\right] \otimes [\Delta p(D1 = \mathsf{ok})]\end{aligned} \quad (4)$$

or as:

$$[\Delta p(L1 = \mathsf{fault})] = \left[\frac{dp(L1 = \mathsf{fault})}{dp(D1 = \mathsf{inst})}\right] \otimes [\Delta p(D1 = \mathsf{inst})] \quad (5)$$

where $\otimes$ and $\oplus$ are qualitative multiplication and addition (Table 5) respectively. Note that in the second equation we could use the derivative with respect to any of the values of D. Note also that the

$\otimes$	[+]	[0]	[−]	[?]
[+]	[+]	[0]	[−]	[?]
[0]	[0]	[0]	[0]	[0]
[−]	[−]	[0]	[+]	[?]
[?]	[?]	[0]	[?]	[?]

$\oplus$	[+]	[0]	[−]	[?]
[+]	[+]	[+]	[?]	[?]
[0]	[+]	[0]	[−]	[?]
[−]	[?]	[−]	[−]	[?]
[?]	[?]	[?]	[?]	[?]

Table 5: Qualitative multiplication and addition

derivatives cannot take arbitrary (qualitative) values, but are constrained to obey the laws of probability. Namely, the normalisation condition (3) given above ensures that the derivative $[\partial p(Li = \mathsf{fault})/\partial p(D)]$ may be calculated as $[-] \otimes [\partial p(Li = \mathsf{ok})/\partial p(D)]$ and $[\partial p(S1 = \mathsf{fault})/\partial p(D)] = [-] \otimes [\partial p(S1 = \mathsf{ok})/\partial p(D)]$ for any value of D. These constraints hold whatever numerical values are used in the model. Thus, irrespective of the way in which a piece of evidence affects the probability of a line or busbar fault, it will influence the probability of the component being ok in the opposite way.

3.3 The Possibility Case

Possibility theory is essentially qualitative [24] when it relies on maximum and minimum operations to determine the updating of values, numbers being used simply as a convenient and easily comprehensible set of suitably ordered values. The use of the max and min operations makes it impossible to obtain true derivatives, but a form of derivative based on small finite, rather than inifintesimal, change may be established [9]. Using the results established there, we can predict that:

$$\left[\frac{\delta\Pi(Li=\mathsf{fault})}{\delta\Pi(D=\mathsf{inst})}\right] = [+] \quad \text{if} \quad \Pi^*(D=\mathsf{inst}) < \Pi(Li=\mathsf{fault} \mid D=\mathsf{inst}) \quad \text{and} \quad \Pi(Li=\mathsf{fault}, D=\mathsf{inst}) \geq \Pi(Li=\mathsf{fault}, D \neq \mathsf{inst}) \tag{6}$$

where $\Pi^*(\cdot)$ is the possibility once some information is received from the alarms[6]. Otherwise the derivative has the value [0]. Similar expressions will relate changes in $\Pi(Li=\mathsf{fault})$ to those in $\Pi(D=\mathsf{del})$ and $\Pi(D=\mathsf{ok})$.

We can also chart the behaviour of the possibility of the line being ok. In (normalized) possibility theory this is only connected to the possibility of line fault by the relation:

$$\max\left(\Pi(Li=\mathsf{ok}), \Pi(Li=\mathsf{fault})\right) = 1 \tag{7}$$

so that one of the two must be perfectly possible. Thus when no fault is observed, it must be perfectly possible that the line is ok, and vice-versa. However, unlike the probability case, in general we are not able to determine how the possibility of the line being ok changes when we know how the possibility of the line being faulty changes. As a result, we have to separately determine how the possibility of a line being ok changes with knowledge of alarms:

$$\left[\frac{\delta\Pi(Li=\mathsf{ok})}{\delta\Pi(D=\mathsf{inst})}\right] = [+] \quad \text{if} \quad \Pi^*(D=\mathsf{inst}) < \Pi(Li=\mathsf{ok} \mid D=\mathsf{inst}) \quad \text{and} \quad \Pi(Li=\mathsf{ok}, D=\mathsf{inst}) \geq \Pi(Li=\mathsf{ok}, D \neq \mathsf{inst}) \tag{8}$$

Again similar expressions will relate changes in $\Pi(Li = \mathsf{ok})$ to those in $\Pi(D = \mathsf{del})$ and $\Pi(D = \mathsf{ok})$. Furthermore, we can relate changes in the possibility of busbar failure, and of the busbar being ok, to instantaneous alarm, delayed alarm, and no alarm. Now, the difference expressions such as $[\delta\Pi(Li=\mathsf{fault})/\delta\Pi(D=\mathsf{inst})]$ that we are using are partial in that they only take account of changes with respect to one possibility. As a result, we calculate the change in possibility of a line fault given information about alarms using:

$$\begin{aligned}[\Delta\Pi(Li=\mathsf{fault})] = {} & \left[\frac{\delta\Pi(Li=\mathsf{fault})}{\delta\Pi(D=\mathsf{inst})}\right] \otimes [\Delta\Pi(D=\mathsf{inst})] \\ & \oplus \left[\frac{\delta\Pi(Li=\mathsf{fault})}{\delta\Pi(D=\mathsf{del})}\right] \otimes [\Delta\Pi(D=\mathsf{del})] \\ & \oplus \left[\frac{\delta\Pi(Li=\mathsf{fault})}{\delta\Pi(D=\mathsf{ok})}\right] \otimes [\Delta\Pi(D=\mathsf{ok})]\end{aligned} \tag{9}$$

and changes in the possibilities of the line being ok and busbar being both ok and faulty may be calculated in an analagous way. As noted above, the normalisation condition only weakly restricts the way possibility values change—it only prevents all the values of some variable from having a possibility of less than one at the same time, and this requirement does not greatly restrict the qualitative derivatives.

[6] The fact that we know how the possibility values change allows us to make more precise predictions than is generally the case [9]

3.4 The Belief Function Case

Once again we can use the results of [9] to determine the qualitative behaviour of the belief function model of the distribution network. Since we are dealing with belief functions, D may take any value in the power set of $\{\mathsf{inst}, \mathsf{del}, \mathsf{ok}\}$, so $D \in \{\mathsf{inst}, \mathsf{del}, \mathsf{ok}, \mathsf{inst} \cup \mathsf{del}, \mathsf{inst} \cup \mathsf{ok}, \mathsf{del} \cup \mathsf{ok}, \mathsf{inst} \cup \mathsf{del} \cup \mathsf{ok}\}$, where we write inst for $\{\mathsf{inst}\}$, del for $\{\mathsf{del}\}$ and ok for $\{\mathsf{ok}\}$ for notational simplicity. As in the probability case, we can either write down partial derivatives of the form:

$$\frac{\partial bel(Li = \mathsf{fault})}{\partial bel(D = \mathsf{inst})} = bel(Li = \mathsf{fault} \mid D = \mathsf{inst}) \tag{10}$$

or more useful total derivatives such as $dbel(Li = \mathsf{fault})/dbel(D = \mathsf{inst})$, which give changes in $bel(Li = \mathsf{fault})$ that also take into account changes that take place in $bel(D = \mathsf{del})$ and $bel(D = \mathsf{ok})$. For the alarm situations that we are interested in, we have, for instance:

$$\begin{aligned} \left[\frac{dbel(Li = \mathsf{fault})}{dbel(D = \mathsf{inst})}\right] = & \\ & [bel(Li \mid D = \mathsf{inst}) - \min_X bel(Li = \mathsf{fault} \mid X)] \\ & \oplus [bel(Li = \mathsf{fault} \mid D = \mathsf{inst}) - \max_X bel(Li = \mathsf{fault} \mid X)] \end{aligned} \tag{11}$$

where $X \subseteq \{D = \mathsf{inst}, D = \mathsf{del}, D = \mathsf{ok}\}$, $D = \mathsf{inst} \in X$. Obviously we have analagous results for other alarm types, and similar results about the busbar may be deduced. As in the probability case the derivatives are not theoretically unconstrained, but since in Dempster-Shafer theory there is only a very weak link between the belief in, say, a line being ok and being faulty:

$$bel(Li = \mathsf{ok}) + bel(Li = \mathsf{fault}) \leq 1 \tag{12}$$

there are, in practice, no restrictions on the qualitative values of the derivatives, indicating that, in qualitative terms, belief functions are a generalisation of both probability and possibility theories.

3.5 More Complex Models

Because our case study is concerned with the analysis of a relatively simple model, in that there is only one arc in the causal graph between any observation that we can make and any hypothesis in which we are interested, it is worth considering how more complex models may be analysed. In fact, the answer is simple. Because the results that we have used are the results of a truly local analysis—that is the analysis of each arc may be carried out using just data about the conditional values that control the behaviour of that arc—we can simply analyse the model arc by arc, and then combine the results along the path from observation to interesting hypothesis.

For example, if we have an observation O, a hypothesis H and an intermediate variable I, so that there is an arc from H to I and from I to O, we would analyse the model by determining $dval(I = i_j)/dval(O = o_i)$ for every value o_i of O and i_j of I, and then $dval(H = h_l)/dval(I = i_j)$ for every value h_l of H and i_j of I. Then to determine the effect of an observation that has the effect of making a given $val(o_i)$ increase, we would calculate the change in a given $val(h_l)$ as:

$$[\Delta val(h_l)] = \left[\frac{dval(H = h_l)}{dval(I = i_j)}\right] \otimes \left[\frac{dval(I = i_j)}{dval(O = o_i)}\right] \otimes [\Delta val(o_i)] \tag{13}$$

where i_j is a suitably convenient value of I. In others words, large models may be handled by decomposing them, analysing the components and combining the results in a modular way.

4 VALIDATING THE BEHAVIOUR OF THE MODELS

Having now analysed the way in which qualitative uncertainty values are propagated through the kind of network structures found in our test case, and thus determined the general behaviour of the formalisms in which we are interested, we can use the numerical information of Section 2 to examine the exact behaviour of the models we have proposed. This provides a qualitative comparison of the different formalisms for our problem, as well as predicting how the quantitative formalisms will behave.

The way new information coming from a circuit breaker is incorporated in the model is different for the three formalisms considered. In probability theory, the information that a particular alarm has arrived from a breaker is typically introduced by increasing the value for the associated state, at the expense of the values of the alternative states: for example, a report of an instantaneous alarm is encoded by forcing $\Delta p(D\!=\!\mathsf{inst}) = [+]$, $\Delta p(D\!=\!\mathsf{del}) = [-]$, and $\Delta p(D\!=\!\mathsf{ok}) = [-]$. In Dempster-Shafer theory, an alarm report is encoded by simply increasing the value of the associated state (the values of altermative states do not need to be decreased as they are initially 0); so, a report of an instantaneous alarm is encoded by $\Delta bel(D\!=\!\mathsf{inst}) = [+]$, $\Delta bel(D\!=\!\mathsf{del}) = [0]$, and $\Delta bel(D\!=\!\mathsf{ok}) = [0]$. Possibility theory differs in that an alarm report is not encoded by increasing the possibility value (which is what we are measuring) of that alarm, but by decreasing the possibility of all other alarm conditions, so that for an instantaneous alarm $\Delta\Pi(D\!=\!\mathsf{inst}) = [0]$, $\Delta\Pi(D\!=\!\mathsf{del}) = [-]$, and $\Delta\Pi(D\!=\!\mathsf{ok}) = [-]$.

4.1 The Probability Case

The analysis of Section 3.2 told us which conditional probabilities determine the behaviour of the system. Now, from Table 1 we can see that for the inner circuit breakers we have:

$$\begin{aligned}
p(Li\!=\!\mathsf{fault}, S1\!=\!\mathsf{fault} \,|\, D1i\!=\!\mathsf{inst}) &= 0.01\\
p(Li\!=\!\mathsf{fault}, S1\!=\!\mathsf{fault} \,|\, D1i\!=\!\mathsf{del}) &= 0.01\\
p(Li\!=\!\mathsf{fault}, S1\!=\!\mathsf{fault} \,|\, D1i\!=\!\mathsf{ok}) &= 0.05\\
p(Li\!=\!\mathsf{fault}, S1\!=\!\mathsf{ok} \,|\, D1i\!=\!\mathsf{inst}) &= 0.89\\
p(Li\!=\!\mathsf{fault}, S1\!=\!\mathsf{ok} \,|\, D1i\!=\!\mathsf{del}) &= 0.89\\
p(Li\!=\!\mathsf{fault}, S1\!=\!\mathsf{ok} \,|\, D1i\!=\!\mathsf{ok}) &= 0.05
\end{aligned}$$

so that $p(Li\!=\!\mathsf{fault} \,|\, D1i\!=\!\mathsf{inst}) = 0.90$, $p(Li\!=\!\mathsf{fault} \,|\, D1i\!=\!\mathsf{del}) = 0.90$, and $p(Li\!=\!\mathsf{fault} \,|\, D1i\!=\!\mathsf{ok}) = 0.01$. Using these values in (2) we find that the qualitative values of the derivatives that link the probability of a line fault to that of an alarm for the inner circuit breakers are:

$$\left[\frac{dp(Li\!=\!\mathsf{fault})}{dp(D1i\!=\!\mathsf{inst})}\right] = [+] \quad \left[\frac{dp(Li\!=\!\mathsf{fault})}{dp(D1i\!=\!\mathsf{del})}\right] = [+] \quad \left[\frac{dp(Li\!=\!\mathsf{fault})}{dp(D1i\!=\!\mathsf{ok})}\right] = [-]$$

Now, from our knowledge of the prior probability values we can see that when we have an instantaneous alarm the change in the probability of a line fault calculated by (5) is:

Report	none	inst	delayed	ok
$[\Delta p(Li\!=\!\mathsf{fault})]$	[0]	[+]	[+]	[−]

In both these cases, the probability of the line being ok, $p(Li\!=\!\mathsf{ok})$ may be calculated from $1 - p(Li\!=\!\mathsf{fault})$, and thus varies inversely to $p(Li\!=\!\mathsf{fault})$. For the outer circuit breakers:

$$\left[\frac{dp(Li\!=\!\mathsf{fault})}{dp(Di\!=\!\mathsf{inst})}\right] = [+] \quad \left[\frac{dp(Li\!=\!\mathsf{fault})}{dp(Di\!=\!\mathsf{del})}\right] = [?] \quad \left[\frac{dp(Li\!=\!\mathsf{fault})}{dp(Di\!=\!\mathsf{ok})}\right] = [-]$$

where the [?] indicates that it is not possible to establish whether the value of the derivative is positive, negative, or zero. Thus we have:

Report	none	inst	delayed	ok
$[\Delta p(Li\!=\!\mathsf{fault})]$	[0]	[+]	[?]	[−]

Here the [?] indicates that we can not predict precisely whether the probability of a line fault will increase, decrease, or not change based on purely qualitative information (the same result will also be obtained from (4) if we use the partial derivatives derived from (1)). However, it is possible to heuristically refine the prediction using some numerical order-of-magnitude information. For instance, taking the quantitative expressions from (1) and the relevant conditional values, comparing the values of the partial derivatives we find that:

$$\frac{\partial p(Li\!=\!\mathsf{fault})}{\partial p(Di\!=\!\mathsf{inst})} > \frac{\partial p(Li\!=\!\mathsf{fault})}{\partial p(Di\!=\!\mathsf{del})} \gg \frac{\partial p(Li\!=\!\mathsf{fault})}{\partial p(Di\!=\!\mathsf{ok})}$$

where $\gg$ indicates a difference of at least an order of magnitude. Now, from the prior probabilities we know that when a delayed alarm takes place $|\Delta p(Di = \mathsf{ok})| \approx |\Delta p(Di = \mathsf{del})| \gg |\Delta p(Di = \mathsf{inst})|$ since the change in $p(Di = \mathsf{ok})$ and $p(Di = \mathsf{del})$ is around a thousand times that in $p(Di = \mathsf{inst})$. Thus when evaluating (4) to establish the change in line fault probability for a delayed alarm, the second term dominates, and we have $[\Delta p(Li\!=\!\mathsf{fault})] = [+]$. Thus, overall:

Report	none	inst	delayed	ok
$[\Delta p(Li\!=\!\mathsf{fault})]$	[0]	[+]	[+]	[−]

Thus the outer circuit breakers work as intended, with the probability of failure of the line increasing with instantaneous and delayed alarms. Thus for both inner and outer breakers, the probability of the line being ok falls with both instantaneous and delayed alarms, and the model is an accurate representation of the behaviour of the target system. Now considering busbar faults; for the inner circuit breakers we have:

$$\left[\frac{dp(S1\!=\!\mathsf{fault})}{dp(D1i\!=\!\mathsf{inst})}\right] = [-] \quad \left[\frac{dp(S1\!=\!\mathsf{fault})}{dp(D1i\!=\!\mathsf{del})}\right] = [-] \quad \left[\frac{dp(S1\!=\!\mathsf{fault})}{dp(D1i\!=\!\mathsf{ok})}\right] = [+]$$

Report	none	inst	delayed	ok
$[\Delta p(S1\!=\!\mathsf{fault})]$	[0]	[−]	[−]	[+]

This means that if we have a report of any kind of alarm in the inner breakers then the probability of a busbar fault decreases, while knowing for sure that there is no alarm means that the probability of failure increases. This behaviour is rather odd since we would expect knowledge of an alarm to increase the probability of a fault, and a no-alarm report to decrease the probability of a fault. We will consider this strange behaviour again in Section 5. Finally, for the outer circuit breakers we have:

$$\left[\frac{dp(S1\!=\!\mathsf{fault})}{dp(Di\!=\!\mathsf{inst})}\right] = [?] \quad \left[\frac{dp(S1\!=\!\mathsf{fault})}{dp(Di\!=\!\mathsf{del})}\right] = [-] \quad \left[\frac{dp(S1\!=\!\mathsf{fault})}{dp(Di\!=\!\mathsf{ok})}\right] = [-]$$

and the purely qualitative approach gives:

Report	none	inst	delayed	ok
$[\Delta p(S1\!=\!\mathsf{fault})]$	[0]	[?]	[+]	[−]

This time order of magnitude considerations cannot help us, and in order to resolve the qualitative ambiguity, we will have to carry out a full numerical calculation using Pulcinella. When this is done, we find, as we would hope, that an instantaneous alarm will cause the probability of a busbar fault to increase. Overall we have:

Report	none	inst	delayed	ok
$[\Delta p(S1\!=\!\mathsf{fault})]$	[0]	[+]	[+]	[−]

which are much as we would expect.

4.2 The Possibility Case

For the case in which the model is quantified using possibility values we find that because the derivatives depend upon the final possibility values of the alarm states, $\Pi^*(D\!=\!\mathsf{inst})$, $\Pi^*(D\!=\!\mathsf{del})$ and $\Pi^*(D\!=\!\mathsf{ok})$ (6) we have different values for different pieces of evidence. For an instantaneous or delayed alarm we have:

$$\left[\frac{\delta\Pi(Li\!=\!\mathsf{fault})}{\delta\Pi(D\!=\!\mathsf{inst})}\right] = [0] \quad \left[\frac{\delta\Pi(Li\!=\!\mathsf{fault})}{\delta\Pi(D\!=\!\mathsf{del})}\right] = [0] \quad \left[\frac{\delta\Pi(Li\!=\!\mathsf{fault})}{\delta\Pi(D\!=\!\mathsf{ok})}\right] = [0]$$

while for no alarm we have:

$$\left[\frac{\delta\Pi(Li\!=\!\mathsf{fault})}{\delta\Pi(D\!=\!\mathsf{inst})}\right] = [+] \quad \left[\frac{\delta\Pi(Li\!=\!\mathsf{fault})}{\delta\Pi(D\!=\!\mathsf{del})}\right] = [+] \quad \left[\frac{\delta\Pi(Li\!=\!\mathsf{fault})}{\delta\Pi(D\!=\!\mathsf{ok})}\right] = [0]$$

so that from (9):

Report	none	inst	delayed	ok
$[\Delta\Pi(Li\!=\!\mathsf{fault})]$	[0]	[0]	[0]	[−]

and the model behaves as expected. If there is an instantaneous or delayed alarm on either the inner or outer breakers, a line fault remains perfectly possible. If there is no alarm the possibility of a line fault falls. We can also consider how the possibility of the line being ok changes. We find that for both types of breaker, for all alarm states:

$$\left[\frac{\delta\Pi(Li\!=\!\mathsf{ok})}{\delta\Pi(D\!=\!\mathsf{inst})}\right] = [0] \quad \left[\frac{\delta\Pi(Li\!=\!\mathsf{ok})}{\delta\Pi(D\!=\!\mathsf{del})}\right] = [0] \quad \left[\frac{\delta\Pi(Li\!=\!\mathsf{ok})}{\delta\Pi(D\!=\!\mathsf{ok})}\right] = [+]$$

which gives:

Report	none	inst	delayed	ok
$[\Delta\Pi(Li\!=\!\mathsf{ok})]$	[0]	[−]	[−]	[0]

So, for both instantaneous and delayed alarms in both inner and outer breakers, the possibility of the line being ok will fall. Thus the possibility of the line being ok does not change[7] as long as there is no alarm on either type of breaker, but falls as soon as we have any kind of alarm. Similar results may be obtained for the possibilities that the busbar has a fault or is ok. It is clear that we can derive conditions similar to those of (6) and (8) to predict how the possibility of the busbar being faulty depends upon the possibilitites of the different types of alarm. We find that for the outer breakers we have:

$$\left[\frac{\delta\Pi(S1\!=\!\mathsf{fault})}{\delta\Pi(Di\!=\!\mathsf{inst})}\right] = [0] \quad \left[\frac{\delta\Pi(S1\!=\!\mathsf{fault})}{\delta\Pi(Di\!=\!\mathsf{del})}\right] = [0] \quad \left[\frac{\delta\Pi(S1\!=\!\mathsf{fault})}{\delta\Pi(Di\!=\!\mathsf{ok})}\right] = [0]$$

for instantaneous and delayed alarms, while no alarm gives us:

$$\left[\frac{\delta\Pi(S1\!=\!\mathsf{fault})}{\delta\Pi(Di\!=\!\mathsf{inst})}\right] = [+] \quad \left[\frac{\delta\Pi(S1\!=\!\mathsf{fault})}{\delta\Pi(Di\!=\!\mathsf{del})}\right] = [+] \quad \left[\frac{\delta\Pi(S1\!=\!\mathsf{fault})}{\delta\Pi(Di\!=\!\mathsf{ok})}\right] = [0]$$

while for the inner breakers we have, for all alarm states:

$$\left[\frac{\delta\Pi(S1\!=\!\mathsf{fault})}{\delta\Pi(D1i\!=\!\mathsf{inst})}\right] = [0] \quad \left[\frac{\delta\Pi(S1\!=\!\mathsf{fault})}{\delta\Pi(D1i\!=\!\mathsf{del})}\right] = [0] \quad \left[\frac{\delta\Pi(S1\!=\!\mathsf{fault})}{\delta\Pi(D1i\!=\!\mathsf{ok})}\right] = [0]$$

Thus from the analagous result to (6), for the outer breakers we get:

[7] Notice that the values of both $\Pi(Li = \mathsf{fault})$ and $\Pi(Li = \mathsf{ok})$ are initially 1. As was the case for the assignment of priors to circuit breaker states, possibility theory (as well as belief functions—see below) allows us to remain agnostic about whether a fault is *a priori* more likely than no fault. This is again not the case for probability theory, due to the additivity axiom.

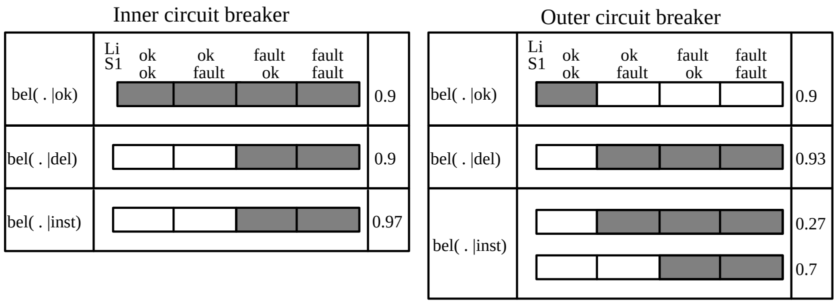

Figure 4: Conditional belief functions for the "alarm" relations

Report	none	inst	delayed	ok
[$\Delta\Pi(S1 = \textsf{fault})$]	[0]	[0]	[0]	[−]

while for the inner breakers, we get:

Report	none	inst	delayed	ok
[$\Delta\Pi(S1 = \textsf{fault})$]	[0]	[0]	[0]	[0]

So, when the possibility of instantaneous and delayed alarms in the outer breakers increases, the possibility of the busbar being faulty will not fall, but this possibility is independent from that of an alarm in the inner breakers. In a similar way, we can determine how the possibility of the busbar being ok will vary. The reader should by now be able to perform all the computations herself, so we only indicate the results—for the outer circuit breakers:

Report	none	inst	delayed	ok
[$\Delta\Pi(S1 = \textsf{ok})$]	[0]	[0]	[−]	[0]

and for the inner circuit breakers:

Report	none	inst	delayed	ok
[$\Delta\Pi(S1 = \textsf{ok})$]	[0]	[0]	[0]	[0]

As we would expect, only the detection of a delayed alarm on the outer breakers suggests that the possibility of the busbar being ok should fall, while no alarm state on the inner breakers will cause the possibility to change. In conclusion, the qualitative behaviour of the model developed using possibility theory corresponds to our expectations given the domain knowledge.

4.3 The Belief Function Case

The belief function case is rather different to the others in that the values given in Section 2 do not correspond to those needed in (10) and (11) to establish the qualitative values of the derivatives. Instead we have to extract the conditional beliefs from the joint mass assignment given in Section 2.2. In order to compute the mass distribution of, say, $L1$ and $S1$ given that $D1 = \textsf{del}$, that is $bel(L1, S1 \,|\, D1 = \textsf{del})$, we just consider all the second rows in the joint distributions (those corresponding to $D1 = \textsf{del}$), and combine them in the usual way for belief functions [4]. Thus, for the outer set of circuit breakers we compute the value of $bel(L1, S1 \,|\, D1 = \textsf{del})$ as $(0.9 \times (0.7 + 0.3) \times 0.7) + (0.9 \times (0.7 + 0.3) \times 0.7) + (0.1 \times (0.7 + 0.3) \times 0.3))$. The full set of conditional assignments over Li and $S1$ is shown in Figure 4, where, for instance, the second row of the "Inner circuit breaker" column gives us $bel(Li, S1 \,|\, D1i = \textsf{del})$ for the various values of $S1$ and any Li. Similarly, the first row of the "Outer circuit breaker" column gives us the values of $bel(Li, S1 \,|\, Di = \textsf{ok})$.

From this we can establish that for the inner circuit breakers $bel(Li = \textsf{ok} \mid D1i = \textsf{ok}) = 0.9$,

$bel(Li = \mathsf{fault} \mid D1i = \mathsf{del}) = 0.9$, and $bel(Li = \mathsf{fault} \mid D1i = \mathsf{inst}) = 0.97$, while for the outer circuit breakers, $bel(Li = \mathsf{ok} \cap S1 = \mathsf{ok} \mid Di = \mathsf{ok}) = 0.9$, $bel(Li = \mathsf{fault} \cup S1 = \mathsf{fault} \mid Di = \mathsf{del}) = 0.93$, $bel(Li = \mathsf{fault} \cup S1 = \mathsf{fault} \mid Di = \mathsf{inst}) = 0.27$, $bel(Li = \mathsf{fault} \mid Di = \mathsf{inst}) = 0.7$ and all other conditional beliefs are zero[8]. From $bel(Li = \mathsf{ok} \cap S1 = \mathsf{ok} \mid Di = \mathsf{ok}) = 0.9$ we know that $bel(Li = \mathsf{ok} \mid Di = \mathsf{ok}) \geq 0.9$ and $bel(S1 = \mathsf{ok} \mid Di = \mathsf{ok}) \geq 0.9$. From (11) we learn that for the inner circuit breakers we have:

$$\left[\frac{dbel(Li=\mathsf{fault})}{dbel(D1i=\mathsf{inst})}\right] = [+] \quad \left[\frac{dbel(Li=\mathsf{fault})}{dbel(D1i=\mathsf{del})}\right] = [+] \quad \left[\frac{dbel(Li=\mathsf{fault})}{dbel(D1i=\mathsf{ok})}\right] = [0]$$

while for the outer circuit breakers:

$$\left[\frac{dbel(Li=\mathsf{fault})}{dbel(Di=\mathsf{inst})}\right] = [+] \quad \left[\frac{dbel(Li=\mathsf{fault})}{dbel(Di=\mathsf{del})}\right] = [0] \quad \left[\frac{dbel(Li=\mathsf{fault})}{dbel(Di=\mathsf{ok})}\right] = [0]$$

When we have evidence of an instantaneous alarm, and belief in an instantaneous alarm thus increases, there is no change in belief in a delayed alarm or no alarm. With this knowledge we can predict that the model behaves as it should. Indeed, for the inner circuit breakers we have:

Report	none	inst	delayed	ok
$[\Delta bel(Li=\mathsf{fault})]$	[0]	[+]	[+]	[0]

while for the outer circuit breakers:

Report	none	inst	delayed	ok
$[\Delta bel(Li=\mathsf{fault})]$	[0]	[+]	[0]	[0]

Thus, if there is an instantaneous or delayed alarm from the inner breakers, belief in a line fault increases. An instantaneous alarm from the outer breakers also increases belief in a line fault, but a delayed alarm does not affect this belief. Knowing that there is no alarm from the breakers does not affect belief in a line fault. These behaviours do not exactly match the specifications in Section 2—namely, we were expecting $bel(Li=\mathsf{fault})$ to increase in response to a delayed alarm in an outer breaker. A similar phenomenon appears in other cases, and we will discuss it in the next section. Using (11) to predict how the our belief in the line being ok changes. We find that for both inner and outer circuit breakers:

$$\left[\frac{dbel(Li=\mathsf{ok})}{dbel(D=\mathsf{inst})}\right] = [0] \quad \left[\frac{dbel(Li=\mathsf{ok})}{dbel(D=\mathsf{del})}\right] = [0] \quad \left[\frac{dbel(Li=\mathsf{ok})}{dbel(D=\mathsf{ok})}\right] = [+]$$

from which we know:

Report	none	inst	delayed	ok
$[\Delta bel(Li=\mathsf{ok})]$	[0]	[0]	[0]	[+]

So, for both instantaneous and delayed alarms in both inner and outer breakers, our belief in the line being ok is unchanged, but this belief will increase when we have evidence that there is no alarm. Similar results may be established for the busbar. Since all the relevant conditional values are zero except $bel(S1 = \mathsf{ok} \mid D = \mathsf{ok})$, the only derivative relevant to the busbar that is non-zero is:

$$\left[\frac{dbel(S1=\mathsf{ok})}{dbel(Di=\mathsf{ok})}\right] = [+]$$

for the outer breakers, meaning that, when we look at how the beliefs change we find that, for the inner breakers:

[8] Note that for notational simplicity we write $bel(Li = \mathsf{ok} \cap S1 = \mathsf{ok} \mid Di = \mathsf{ok})$ to denote $bel(\langle Li, S1\rangle = \{\langle \mathsf{ok}, \mathsf{ok}\rangle\} \mid Di = \mathsf{ok})$, the conditional belief that the joint variable $\langle Li, S1\rangle$ takes the value $\langle \mathsf{ok}, \mathsf{ok}\rangle$. $bel(Li = \mathsf{fault} \cup S1 = \mathsf{fault} \mid Di = \mathsf{inst})$ is used in a similar way.

Report	none	inst	delayed	ok
[$\Delta bel(S1=\mathsf{fault})$]	[0]	[0]	[0]	[0]
[$\Delta bel(S1=\mathsf{ok})$]	[0]	[0]	[0]	[0]

while for the outer breakers:

Report	none	inst	delayed	ok
[$\Delta bel(S1=\mathsf{fault})$]	[0]	[0]	[0]	[0]
[$\Delta bel(S1=\mathsf{ok})$]	[0]	[0]	[0]	[+]

Thus, belief in the busbar failing or being ok is insensitive to anything happening on the inner breakers, as well as to instantaneous and delayed faults on the outer breaker. Belief in the busbar being ok increases when it is known that there is no fault on the outer breaker which is the only one that points to the busbar. Overall, the belief function model of the alarm system behaves largely as one would expect on reading the description in Section 2 with the exception of the anomaly described above—belief in a fault in the busbar does not increase with belief in a delayed alarm in an outer breaker.

4.4 Discussion of Verification

The main result that emerges from this section is that we have shown how the qualitative analysis may be used to validate the quantitative model. It can do this since it is possible to make qualitative predictions of the behaviour of the quantitative model using the numerical values. This qualitative prediction may be compared against the original opinion of the domain expert to determine whether the quantitative model has captured the expert's knowledge. We emphasize that this technique is analytic rather than experimental, and can be carried out for formalisms other than probability theory. The usefulness of our technique is highlighted by the fact that it threw up some anomalies in the quantitative model that we were using (these are examined in the next section). In other words, the initial valuation system model, with values as ascertained by the knowledge engineer, does not behave quite as might be expected from the description of its intended behaviour that is supplied in Section 2. Since the validation is purely qualitative, the extent of any anomalies cannot be predicted without performing a numerical calculation. The fact that the validation is purely qualitative also means that there are circumstances under which it gives ambiguous results. Sometimes these can be resolved using a form of order-of-magnitude reasoning (which is formalised in [25]), and sometimes it is necessary to carry out a numerical calculation to tell exactly what happens.

5 DEBUGGING THE VALUES

The qualitative analysis of the probabilistic model of the distribution network revealed that any alarm on an inner circuit breaker will cause the probability of busbar failure to decrease, whilst the observation of "no alarm" will actually cause the probability of busbar failure to increase! This kind of behaviour was a considerable surprise to the knowledge engineer who performed the original elicitation, and it was certainly not intended when the original model was built. Furthermore, because it was so unexpected, such behaviour is unlikely to have been exposed without the qualitative analysis. Now that the anomaly has been spotted, we can run Pulcinella over the specific data, and the quantitative data can be examined to evaluate the impact of the discrepancy. In this case we find that:

Report	none	inst	delayed	ok
$p(S1=\mathsf{fault})$	0.000175	0.000088	0.000088	0.000176

This makes it clear that the change of value is so small that it is probably not bothering with. However, even though the values are small, the fact that the behaviour is reversed for all the possible states of the breaker may induce us to try to debug it.

A second anomaly was detected in the belief function model. This occurs in the system's response to delayed alarms from the outer breakers—such an alarm should, according to the specifications in Section 2.1, increase belief in both a busbar fault and in a fault in the line that the breaker is on. However, in the model no such effects are observed, with the alarm failing to change any of these beliefs. In fact, running Pulcinella on this scenario yields the following results for the outer circuit breakers:

Report	none	inst	delayed	ok
$bel(Li{=}\mathsf{fault})$	0	0.98	0	0

After some thought, however, what could have been seen as a bug in the model turns out to be acceptable behaviour. In the case of a delayed alarm in an outer breaker, say $D1$, the belief function model makes the belief in the event "fault in $L1$ OR $S1$" increase; however, this belief is not committed to either of the two fault events individually[9]. This should not to be seen as a loss of inferential power—the arrival of subsequent items of evidence (i.e., the receipt of new alarm reports) will disambiguate the situation. To see this, suppose that a new delayed alarm is received from $D3$. This supports the hypothesis "fault in $L3$ OR $S1$", and the two pieces of evidence together support the hypothesis "fault in $S1$".

So, for both these anomalies it is possible to argue that what has been detected is not really a problem. However, we might wish to "correct" such behaviours so that inner circuit breaker alarms cause the probability of busbar failure to increase (on the grounds that there is an error somewhere) and delayed alarms on a single outer circuit breaker cause belief in busbar and line failure to increase. If we do so, we can use the qualitative analysis to guide us.

5.1 A Procedure for Debugging

The qualitative analysis tells us two things. Firstly it tells us how particular pieces of evidence affect the fault hypotheses and thus enables us to detect when the model deviates from its intended behaviour. It also tells us which conditional values determine the qualitative behaviour, and this is the key to debugging the model. We can use the analysis to identify which conditional values cause the unwanted behaviour, which ones help to determine the correct behaviour, and which ones do not affect either. We can then alter the values that cause the unwanted behaviour but which do not affect the desired behaviour. To do this, the following informal procedure is suggested:

1. Establish the set of incorrect derivatives by comparing the behaviour of the model to the system behaviour given by the domain expert.

2. For every incorrect derivative, identify the set of equations that determine it, and for every equation e construct the corresponding set of conditionals $C(e)$ that are mentioned in it.

3. For each set $C(e)$ so constructed, let $U(e)$ be the set of conditionals in $C(e)$ that are unused by all other equations.

4. If $U(e)$ is non-empty, then find a new distribution of values to the conditionals in it so that the correct value for the corresponding derivatives are obtained, and the new distribution is as close as possible to the old one.

5. If $U(e)$ is empty, then let $U(e)$ be the set of conditionals in $C(e)$ that do not determine any correct derivatives.

[9] As may be seen from the fact that the relevant conditional belief (identified by rewriting (11) with $L1{=}\mathsf{fault} \cup S1{=}\mathsf{fault}$ in place of $L1{=}\mathsf{fault}$), $bel(L1{=}\mathsf{fault} \cup S1{=}\mathsf{fault} \mid D{=}\mathsf{del}) = 0.93$ while $bel(L1{=}\mathsf{fault} \mid D{=}\mathsf{del}) = bel(S1{=}\mathsf{fault} \mid D{=}\mathsf{del}) = 0$.

6. Find a new distribution for the values to the conditionals in $U(e)$ such that every equation which previously determined an incorrect derivative now produces a correct derivative, no derivatives corrected in step 4 are made incorrect, and the new distribution is as close as possible to the old one.

7. If there are certain equations which may not be corrected in this way (because for instance there are no conditionals in step 5 that do not help to determine initially correct derivatives) then identify the set of all such equations.

8. Find a new distribution over the full set of conditionals involved in these equations such that all of them now predict the correct values of the derivatives, no previously correct or corrected derivatives are made incorrect, and the new distribution is as close as possible to the old one.

"As close as possible" should be interpreted relative to the judgement of the knowledge engineer, who decides how large a deviation from intended behaviour can be tolerated, and which deviations are less tolerable than others. Of course the new distribution must conform to the normalisation condition for the calculus the values are expressed in.

5.2 Correcting the Behaviour of the Probabilistic Model

We can use this procedure to correct the behaviour of the probabilistic model regarding the detection of busbar faults on the inner breakers. Section 4.1 gives us the following set of incorrect derivatives since every single behaviour of the model is qualitatively incorrect:

$$\left\{ \left[\frac{dp(S1{=}\mathsf{fault})}{dp(D1i{=}\mathsf{inst})}\right] \quad \left[\frac{dp(S1{=}\mathsf{fault})}{dp(D1i{=}\mathsf{del})}\right] \quad \left[\frac{dp(S1{=}\mathsf{fault})}{dp(D1i{=}\mathsf{ok})}\right] \right\}$$

Now, these are determined by the equations:

$$\left[\frac{dp(S1{=}\mathsf{fault})}{dp(D1i{=}\mathsf{inst})}\right] = [p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{inst}) - p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{del})] \oplus [p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{inst}) - p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{ok})] \tag{14}$$

$$\left[\frac{dp(S1{=}\mathsf{fault})}{dp(D1i{=}\mathsf{del})}\right] = [p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{del}) - p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{inst})] \oplus [p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{del}) - p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{ok})] \tag{15}$$

$$\left[\frac{dp(S1{=}\mathsf{fault})}{dp(D1i{=}\mathsf{ok})}\right] = [p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{ok}) - p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{inst})] \oplus [p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{ok}) - p(S1{=}\mathsf{fault} \mid D1i{=}\mathsf{del})] \tag{16}$$

now, since $p(S1{=}\mathsf{fault} \mid D1i) = p(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \mid D1i) + p(S1{=}\mathsf{fault}, Li{=}\mathsf{ok} \mid D1i)$ for any value of $D1i$, the set of conditionals used by (14), (15), and (16) is (it is the same for all of them):

$$\{p(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \mid D1i{=}\mathsf{inst}), p(S1{=}\mathsf{fault}, Li{=}\mathsf{ok} \mid D1i{=}\mathsf{inst}),\ p(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \mid D1i{=}\mathsf{del}), p(S1{=}\mathsf{fault}, Li{=}\mathsf{ok} \mid D1i{=}\mathsf{del}),\ p(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \mid D1i{=}\mathsf{ok}), p(S1{=}\mathsf{fault}, Li{=}\mathsf{ok} \mid D1i{=}\mathsf{ok})\}$$

Clearly all the conditionals in each set are used in other equations, so we have to look for conditionals that do not help to determine currently correct derivatives. Since conditionals such as $p(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \mid D1i)$ help to determine the currently correct derivatives that relate $p(Li{=}\mathsf{fault})$ to $p(D1i)$, the conditionals we are looking for are:

$$\{p(S1{=}\mathsf{fault}, Li{=}\mathsf{ok} \,|\, D1i{=}\mathsf{inst}) \;, \\ p(S1{=}\mathsf{fault}, Li{=}\mathsf{ok} \,|\, D1i{=}\mathsf{del}), \\ p(S1{=}\mathsf{fault}, Li{=}\mathsf{ok} \,|\, D1i{=}\mathsf{ok})\}$$

We then choose a new distribution to minimise change, and a brief experiment should suffice to satisfy the reader that a suitable new distribution is:

$$\begin{aligned} p(S1{=}\mathsf{fault}, Li{=}\mathsf{ok} \,|\, D{=}\mathsf{ok}) &= 0.15 \\ p(S1{=}\mathsf{fault}, Li{=}\mathsf{ok} \,|\, D{=}\mathsf{inst}) &= 0.2 \\ p(S1{=}\mathsf{fault}, Li{=}\mathsf{ok} \,|\, D{=}\mathsf{del}) &= 0.2 \end{aligned}$$

5.3 Correcting the Behaviour of the Belief Function Model

In the belief function model, the problem is that belief in line fault and busbar fault does not increase when there is a delayed alarm in the outer breakers, in other words, $[dbel(Li{=}\mathsf{fault}, S1{=}\mathsf{fault})/dbel(Di{=}\mathsf{del})] = [0]$ instead of $[+]$. Thus the incorrect derivative is:

$$\left\{ \left[\frac{dbel(S1{=}\mathsf{fault}, Li{=}\mathsf{fault})}{dbel(Di{=}\mathsf{del})} \right] \right\}$$

And so from (11) we can tell, by analogy, that the relevant equation is:

$$\begin{aligned} \left[\frac{dbel(S1{=}\mathsf{fault}, Li{=}\mathsf{fault})}{dbel(Di{=}\mathsf{del})} \right] = \; & [bel(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \,|\, Di{=}\mathsf{inst}) \\ & \quad - \min_X bel(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \,|\, X)] \\ & \oplus [bel(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \,|\, Di{=}\mathsf{inst}) \\ & \quad - \max_X bel(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \,|\, X)] \end{aligned}$$

where $X \subseteq \{Di{=}\mathsf{inst}, Di{=}\mathsf{del}, Di{=}\mathsf{ok}\}$, $Di{=}\mathsf{del} \in X$, so that the set of conditionals in which we are interested is:

$$\{bel(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \,|\, Di{=}\mathsf{del} \cup Di{=}\mathsf{inst} \cup Di{=}\mathsf{ok}) \;, \\ bel(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \,|\, Di{=}\mathsf{del} \cup Di{=}\mathsf{inst}), \\ bel(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \,|\, Di{=}\mathsf{del} \cup Di{=}\mathsf{ok}), \\ bel(S1{=}\mathsf{fault}, Li{=}\mathsf{fault} \,|\, Di{=}\mathsf{del})\}$$

None of these conditional values is used in any other equation, so that any of the conditional values may be changed without altering other derivatives. In addition, all the relevant conditionals are currently zero, so that $[dbel(Li = \mathsf{fault}, S1 = \mathsf{fault})/dbel(Di = \mathsf{del})]$ may be made positive by making $bel(Li = \mathsf{fault}, S1 = \mathsf{fault} \mid Di = \mathsf{del})$ positive while leaving all other relevant conditionals unchanged. This may be achieved by adding another belief function to the basic probability assignment of Section 2.2 so that we have the basic belief assignments of Figure 5. Note that the change in distribution here is not strictly minimal, since setting the mass of the new basic probability assignment to 0.0000...01 would have sufficed to make the relevant derivative positive. However, it seemed reasonable to our knowledge engineer to change the distribution to a value that generates a significant change in $bel(Li{=}\mathsf{fault}, S1{=}\mathsf{fault})$ given the evidence of a delayed alarm.

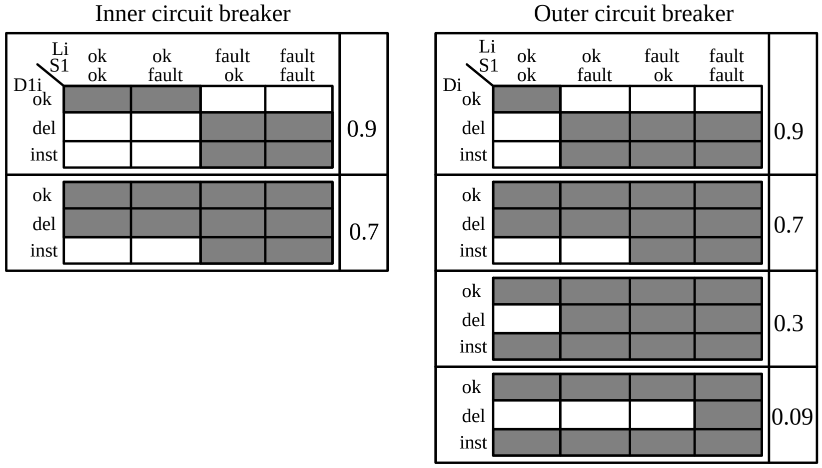

Figure 5: The new joint belief distribution

5.4 System Behaviour after Debugging

Having made the changes outlined in Sections 5.2 and 5.3 we can test the behaviour of the new models. First consider the probabilistic model. Entering the new conditional values into Pulcinella and propagating we obtain, for alarms in the inner circuit breaker:

Report	none	inst	delayed	ok
$p(S1{=}\mathsf{fault})$	0.000007	0.00004	0.0006	0.000006

So an alarm now causes the probability of busbar failure to increase, and the system has had the "bug" removed. Similarly for the belief function model, we have, for the outer breaker:

Report	none	inst	delayed	ok
$bel(S1{=}\mathsf{fault})$	0	0	0.09	0
$bel(Li{=}\mathsf{fault})$	0	0	0.09	0

which corrects the "bug" detected in the original behaviour by relating a delayed alarm from an outer breaker with an increase in the belief in a fault in the busbar and line. Notice that the behaviour for all the other cases remains unchanged.

6 SUMMARY

This paper has investigated one of the uses of reasoning about qualitative changes in uncertainty value in the context of a real world problem. We have illustrated our approach to the qualitative analysis of quantitative models of uncertainty [9] by analysing the behaviour of three models that had been built to diagnose faults in an electricity distribution network. As the analysis is conducted on a local basis, the proposed approach should scale well to larger problems—we can simply analyse the model arc by arc, and then combine the results along the path from observation to interesting hypothesis. The number of steps in the analysis is only dependent upon the number of arcs in the network, and the computational effort to calculate their effect is quadratic in the number of nodes [26]. This analysis is useful because it provides a way of determining the behaviour of a model at a high level of abstraction. As such it relies on weak information and produces results that, although correct, may at times be too weak to be useful. As a result we enriched the purely qualitative analysis by introducing some informal order of magnitude reasoning which gave stronger results. It should be noted that this form of reasoning, like any other kind of order of magnitude reasoning, is essentially heuristic, and may at times cause errors. However, such problems can often be avoided, especially if the order of magnitude reasoning is formalised [25]. Of course, it is also possible to resolve overly weak results in other ways—most obviously by running numerical simulations in the relevant places.

Having performed the analysis, we then used the results to predict the behaviour of models built using all three uncertainty formalisms. This made it easy to check the behaviour of the models against the behaviour which the knowledge engineer who originally elicited the information intended them to have. In particular we checked the way in which the value of a fault hypothesis changed as single pieces of evidence were observed. As a result of this validation a couple of anomalies in behaviour were discovered in which the value of various fault hypotheses did not change as expected. After a discussion of the nature of these anomalies, we presented a procedure for removing them if desired. This procedure makes use of the qualitative analysis to identify which parts of the numerical model must be altered in order to correct the behaviours, and which parts should not be altered in order to prevent changing behaviours that are already correct. We then applied this procedure to correct the anomalies, and were able to demonstrate that after correction the models behaved as they were originally intended to do.

Acknowledgements

The research described in this paper was partially supported by Esprit Projects 3085 DRUMS, 6156 DRUMS II, and 2256 ARCHON. The first author benefitted from a grant from BT, and the second from a grant from the Communauté Française de Belgique under the A.R.C. BELON. Inaki Laresgoiti helped analyze the electricity network problem, and Moisès Goldszmidt identified some problems with the causal network. Both authors would like to thank all their colleagues on DRUMS and DRUMS II for their help, especially Mirko Dohnal, John Fox, Paul Krause, Rudolph Kruse, Philippe Smets, Elisabeth Umkehrer and Hong Xu. In particular, Elisabeth and Hong provided tremendous help in developing and using Pulcinella, while Paul gave a creditable impression of a Bayesian. Discussions with Simon Parsons and Alessandro Saffiotti have also been very useful.

References

[1] Pearl, J. *Probabilistic reasoning in intelligent systems: networks of plausible inference*, Morgan Kaufman, San Mateo CA, 1988.

[2] Zadeh, L. A. Fuzzy sets as a basis for a theory of possibility, *Fuzzy Sets and Systems*, 1, 1–28, 1978.

[3] Dubois, D. and Prade, H. *Possibility theory—an approach to computerized processing of uncertainty*, Plenum Press, New York, 1988.

[4] Shafer, G. *A mathematical theory of evidence*, Princeton University Press, Princeton, NJ, 1976.

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

[6] Henrion, M. and Druzdzel, M. J. Qualitative propagation and scenario-based approaches to explanation of probabilistic reasoning, *Proceedings of the 6th Conference on Uncertainty in Artificial Intelligence*, Boston, Massachusetts, 10–20, 1990.

[7] Wellman, M. P. and Henrion, M. Qualitative Intercausal Relations, or Explaining "Explaining Away", *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, Cambridge, Massachusetts, 535–546, 1991.

[8] Druzdzel, M. and Henrion, M. Belief propagation in qualitative probabilistic networks, in *Qualitative reasoning and decision technologies* (N. Piera Carreté and M. G. Singh, eds.) CIMNE, Barcelona, 451–460, 1993.

[9] Parsons, S. *Qualitative approaches to reasoning under uncertainty*, MIT Press, to appear.

[10] Parsons, S. and Mamdani E. H. On reasoning in networks with qualitative uncertainty, *Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence*, Washington D.C., 435–460, 1993.

[11] Szolovits, P. and Pauker, S. G. Categorical and probabilistic reasoning in medical diagnosis, *Artificial Intelligence*, **11**, 115–144, 1978.

[12] Fox, J. Three arguments for extending the framework of probability, in *Uncertainty in Artificial Intelligence* (L. N. Kanal and J. F. Lemmer eds). North Holland, Amsterdam, 447–458, 1986.

[13] Saffiotti, A. An AI view of the treatment of uncertainty. *The Knowledge Engineering Review*, **2**, 75–97, 1987.

[14] Krause, P. J. and Clark, D. A. *Representing uncertain knowledge—an artificial intelligence approach*, Intellect, 1993.

[15] Neapolitan, R. The interpretation and application of belief functions, *Applied Artificial Intelligence*, **7**, 195–204, 1993.

[16] Shenoy, P. P. and Shafer, G. Axioms for probability and belief function propagation, in: *Uncertainty in Artificial Intelligence 4* (R. D. Shachter, T. S. Levitt, L. N. Kanal, and J. F. Lemmer, eds.). North-Holland, Amsterdam, 169–198, 1990.

[17] Shenoy, P. P. Valuation Based Systems: A Framework for Managing Uncertainty in Expert Systems, University of Kansas, School of Business, Working Paper 226, 1991.

[18] Saffiotti, A., and Umkehrer, E. Pulcinella User's Manual, Technical Report TR/IRIDIA/91-5, Université Libre de Bruxelles, Bruxelles, Belgium, 1991.

[19] Saffiotti, A. and Umkehrer, E. Pulcinella: a tool for propagating uncertainty by local computation. *Proceedings of the 6th Conference on Uncertainty in Artificial Intelligence*, Los Angeles, California, 323–331, 1991.

[20] Saffiotti, A., Parsons, S. and Umkehrer, E. A Case Study in Comparing Uncertainty Management Techniques. *Microcomputers in Civil Engineering, Special Issue on Uncertainty in Expert Systems*, **9**, 367–380, 1994.

[21] Fonck, P. and Straszecka, E. Building influence networks in the framework of possibility theory, *Annales Univ. Sci. Budapest. Sect. Comp.*, **12**, 101–106, 1991.

[22] Smets, Ph. Belief functions: the disjunctive rule of combination and the generalized bayesian theorem, *International Journal of Approximate Reasoning*, **9**, 1–35, 1993.

[23] Bobrow, D. G. *Qualitative reasoning about physical systems*, North-Holland, Amsterdam, 1984.

[24] Dubois, D. Belief structures, possibility theory, decomposable confidence measures on finite sets, *Computers and AI*, **5**, 403–417, 1986.

[25] Parsons, S. Order of magnitude reasoning and qualitative probability, Technical Report, Department of Electronic Engineering, Queen Mary and Westfield College, London.

[26] Druzdzel, M. J. and Henrion, M. "Efficient reasoning in qualitative probabilistic networks", *Proceedings of the 11th National Conference on Artificial Intelligence*, Washington D.C., 548–553, 1993.