Causal reasoning under uncertainty with Q/C-E networks: A case study on preventive diagnosis of power transformers

Pietro Baroni°, Giovanni Guida°, Silvano Mussi*
°*Dipartimento di Elettronica per l'Automazione, Università di Brescia, Brescia, Italy*
**CILEA, Milano, Italy*

Abstract

The paper presents the formalism of quantified causal-evidential networks (Q/C-E networks) for causal reasoning under uncertainty in networks of propositions. First, the basic concept of C-E network is introduced. The issue of representing uncertainty about propositions and causal-evidential relations is then discussed and Q/C-E networks are defined. Methods for propagating and aggregating uncertainty in a Q/C-E network are proposed and their main properties are illustrated. The proposed approach has been successfully experimented in the design and development of ASTRA, a knowledge-based system for preventive diagnosis of power transformers.

1 Introduction

Reasoning under uncertainty in networks of propositions is a key issue in applications of artificial intelligence to diagnosis. Artificial intelligence literature offers a wide range of proposals concerning networks of propositions affected by uncertainty, such as INFERNO [11], belief networks [8], and HUGIN [1]. Bonissone [4] has specified a set of desiderata for an uncertainty management system. Many classical approaches, however, fail to match very important requirements. For instance, the Bayesian approach [9] does not clearly distinguish between ignorance and contradiction, and Dempster-Shafer theory [12] requires a set of mutually exclusive and exhaustive hypotheses to be available. Numerical representation of uncertainty suffers form scarce cognitive plausibility [6]. Driankov's uncertainty calculus [7] eliminates many of these drawbacks; however it is affected by some cognitive inconsistencies, as pointed out, for example, in [3]. This paper presents Quantified Causal-Evidential networks (Q/C-E networks), a framework for reasoning under uncertainty in the context of networks of propositions. This approach integrates the representation of cause-effect and evidence-justification relations between propositions with the management of uncertainty that can affect both propositions and the relations between them. Although some basic concepts of our proposal are shared with Pearl’s belief networks [8] and with Driankov’s uncertainty theory [7], Q/C-E networks are not an extension of any existing knowledge representation technique, but, rather, an original proposal. Q/C-E networks have been experimented in the development of ASTRA, a knowledge-based system for state assessment and preventive diagnosis of power transformers [2].

2 C-E networks

A *Causal-Evidential network* (*C-E network*) is made up of a finite set of *nodes*, representing propositions, and a finite set of *links* between pair of nodes, representing relations between propositions. Each link is made up of two directed arcs: one represents a causal relation and is called a *C-link*, whereas the other, directed in the opposite direction, represents an evidential relation and is called an *E-link*. Thus, if a proposition P is linked with a C-link to a proposition Q (i.e., P $\xrightarrow[C]{}$ Q), then Q is linked with an E-link to P (i.e., Q $\xrightarrow[E]{}$ P). The concept of C-E network is better clarified by means of an example. Let us consider the following two propositions (extracted from ASTRA): "Oil dielectric strength is low" (OIL-DIEL-LOW), "Partial discharges are present" (PARTIAL-DIS). In the model of a power transformer, OIL-DIEL-LOW represents a physical state which plays the role of a "cause" with respect to the physical state represented by PARTIAL-DIS, which plays the role of an "effect". At the same time, if the physical state PARTIAL-DIS is observed, this can be interpreted as an "evidence" which can find a "justification" in the physical state OIL-DIEL-LOW. Therefore, this piece of knowledge can be appropriately represented by the fragment of C-E network shown in Figure 1.

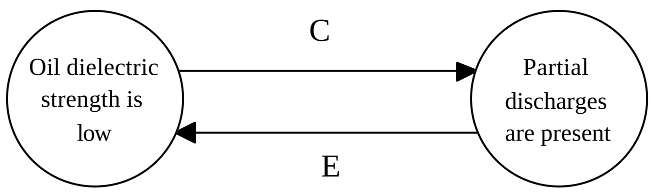

Figure 1: A fragment of C-E network.

3 From C-E networks to Q/C-E networks

3.1 From physical states to mental models

An important aspect of reasoning with C-E networks concerns the possibility of taking into account how much a proposition is believed and how much this belief affects the belief in other propositions connected through C-links or E-links.

Therefore, in order to make the C-E network formalism appropriate to model (some aspects of) human reasoning, the need of quantifying beliefs arises. Starting from these considerations, let us define the concept of quantified C-E network. A *Quantified C-E network* (*Q/C-E network*) is a C-E network where nodes represent beliefs in propositions and links between pairs of nodes represent relations between the corresponding beliefs in propositions.

Whereas propositions of C-E networks represent physical states of the world, propositions of Q/C-E networks, represent beliefs; i.e. mental states of a person who observes the world and reasons about it (for example, a domain expert).

In order to quantify how much a proposition is believed and how much this belief affects the beliefs in other propositions connected to it, we must define the concepts of "quantified proposition" and "quantified relation".

3.2 Quantified propositions

For the sake of clarity let us refer to the proposition "Partial discharges are present" (PARTIAL-DIS). In real cases, such a proposition is not considered by the expert as being definitely true or false, since there are no means of directly verifying it. In fact the expert, on the basis of the available evidence in favor of PARTIAL-DIS, is prompted to believe to a certain extent that PARTIAL-DIS is true, whereas, on the basis of other available evidence against PARTIAL-DIS, the expert may be prompted to believe to a certain extent that PARTIAL-DIS is false. Beliefs may assume values in a finite, ordered set of *belief degrees*. We call *quantified proposition* a proposition P having associated a belief pair [bel(P, true), bel(P, false)], where bel(P, true) and bel(P, false) represent the belief degrees related to the truth and to the falsity of P respectively. The pair [bel(P, true), bel(P, false)] is called the *belief state* of P.

3.3 Quantified relations

Even relations between propositions may be affected by uncertainty. The expert may be not completely sure about the existence or non-existence of a link between any pair of propositions, because of incomplete knowledge about the world. For example, considering again the fragment of C-E network shown in Figure 1, if he is certain that "Oil dielectric strength is low"(OIL-DIEL-LOW), then he may believe to a certain extent (but not completely) that, as a consequence, "Partial discharges are present" (PARTIAL-DIS), in fact there are other (not completely known) factors which may enable or prevent low oil dielectric strength to produce partial discharges.

Therefore, uncertainty can relate both to propositions and to relations between propositions. However, the following distinction should be noticed. Belief degrees associated to a proposition may be revisited whenever new evidence becomes available. The belief state of a quantified proposition is determined on the basis of the knowledge available about the current state of the world, so in a sense it can be considered as being a sort of "fresh knowledge". For example, how much the expert believes that "water concentration in oil is high" (WATER-HIGH) depends on the result of a measure on the current oil sample. On the contrary, how much an expert believes in a relation between two propositions may be considered as being a sort of "consolidated knowledge", resulting from the synthesis of several experiences occurred in the past. It was during these experiences that the expert collected evidences in favour or against a relation between propositions, but the precise record of the past experiences from which this relational knowledge was originally derived is generally lost (it is the so called "paradox of expertise": the expert does not know why he knows). Moreover, relational knowledge only concerns believed relations (i.e. believed at least to a given degree), since relations which are not believed (either disbelieved or unknown) are simply not asserted. In conclusion, uncertainty of relations is represented through the quantification of the extent to which one is authorized to believe in the existence and validity of a relation. Given this premise let us associate a single belief degree to a relation representing how much that relation is believed by the expert. The value of this belief degree may be elicited asking the expert how much he is prompted to believe in Q, supposed that P is totally believed to be true.

We call *quantified relation* a relation $P \longrightarrow Q$ having associated a belief bel($P \longrightarrow Q$, true) that represent the belief degree related to the truth of $P \longrightarrow Q$. The belief bel($P \longrightarrow Q$, true) is called the *belief level* of $P \longrightarrow Q$.

It should be noted that the belief levels of the two opposite links (causal and evidential) connecting a pair of propositions (namely: $P \xrightarrow[C]{} Q$ and $Q \xrightarrow[E]{} P$) are independent quantifications.

3.4 Representing uncertainty with Q/C-E Networks

A belief degree is usually represented as a real number in the closed interval [0, 1]. However, as it is well known, it is not natural for an expert to express a belief judgement in numerical form. Nor it is easy for a user to attach a correct intuitive meaning to a numerical belief. It is much easier and more natural to use vague *linguistic labels*. For example, let us consider the question: if you are sure that "Oil dielectric strength is low" (OIL-DIEL-LOW), what is your degree of belief that "Partial discharges are present" (PARTIAL-DIS)? The expert finds more natural answering by choosing among a set of vague terms like for example {very-low, low, ..., high, very-high}, than choosing a number in [0, 1]. Moreover vague terms may be represented as fuzzy numbers [13], matching this way the requirements of natural knowledge elicitation and presentation with the need of formal calculus. In the case of ASTRA, for example, a scale of nine linguistic labels have been adopted, derived from [5], namely: [UNINFORMED, EXTREMELY UNLIKELY, MOST UNLIKELY, UNLIKELY, IT MAY, LIKELY, MOST LIKELY, EXTREMELY LIKELY, CERTAIN]. Each label in this set is represented in short form as E_i, where i = 1, 2, ..., 9. E_1 (UNINFORMED, the minimum) represents negligible belief (including null belief) in the truth or falsity of a proposition P, while E_9 (CERTAIN, the maximum) represents full belief in the truth or falsity of P.

4 Reasoning under uncertainty in Q/C-E Networks

Reasoning under uncertainty in Q/C-E networks concerns both the propagation of belief states through C- and E-links and the aggregation of belief states concerning the same proposition and originated from different information sources. For the sake of clarity, we examine propagation in two stages: local propagation (i.e., from-one node to the adjacent nodes) and global propagation (i.e., propagation of belief states through the entire Q/C-E network).

4.1 Local propagation

Given two linked propositions P and Q and given the C-link belief levels of the C- and E-link between them, propagating a belief state from P to Q means establishing which belief state should be associated to Q given the belief state associated to P. More precisely, let us define *local propagation* as the problem of deriving [bel(Q, $\underline{\text{true}}$), bel(Q, $\underline{\text{false}}$)] from [bel(P, $\underline{\text{true}}$), bel(P, $\underline{\text{false}}$)], bel(P $\xrightarrow[C]{}$ Q), and bel(Q $\xrightarrow[E]{}$ P).

As far as bel(Q, $\underline{\text{true}}$) is concerned, it is defined as follows:

$$\text{bel}(Q, \underline{\text{true}}) = \text{bel}(P, \underline{\text{true}}) * \text{bel}(P \xrightarrow[C]{} Q),$$

where the symbol $*$ denotes for a T-norm operator (see [4] for a formal definition) between the fuzzy numbers representing belief degrees. This way of calculating bel(Q, $\underline{\text{true}}$) is based on mere common-sense and proves to be sound. In fact, on the basis of the definition of belief level, the computation of bel(Q, $\underline{\text{true}}$) follows from the intuitive proportion:

$$E_9 : \text{bel}(P \xrightarrow[C]{} Q) = \text{bel}(P, \underline{\text{true}}) : x.$$

As far as bel(Q, $\underline{\text{false}}$) is concerned, it should be derived from bel(P, $\underline{\text{false}}$) through bel($\neg$ P $\xrightarrow[C]{}$ $\neg$ Q). However knowledge about $\neg$ P $\xrightarrow[C]{}$ $\neg$ Q is not available in a Q/C-E network. What is available is bel(Q $\xrightarrow[E]{}$ P), but how can this be used to derive bel(Q, $\underline{\text{false}}$)? To answer this question, let us note that, intuitively, the more one is prompted to believe in P, because of belief in Q, the more he is prompted to believe in $\neg$ Q, because of belief in $\neg$ P. This point was confirmed during knowledge elicitation for ASTRA, both in case of an effect with only one cause and in case of an effect with multiple possible causes. For example, as far as the case of an effect with only one cause is concerned, let us consider the following case. When a transformer is in service, the event "Winding buckling"(WIND-B) may only be caused by "Over-current events" (OCE). In this case the belief level at which the expert is prompted to believe in the hypothesis OCE, because of his belief state in WIND-B is very high, i.e. the belief degree of WIND-B $\xrightarrow[E]{}$ OCE is very high. On the other hand, the belief level at which the belief in the absence of over-current events (i.e., $\neg$ OCE)

produces the belief that there is no winding buckling (i.e., $\neg$ WIND-B) is very high too, due to the strong correlation between them.
On the basis of these remarks, we postulate the numerical equality between bel($\neg$ P $\xrightarrow[C]{}$ $\neg$ Q) and bel(Q $\xrightarrow[E]{}$ P). As a consequence bel(Q, false) is defined as follows:

bel(Q, false) = bel(P, false) $*$ bel(Q $\xrightarrow[E]{}$ P).

The discussion of local propagation presented above has considered the case of deriving the belief state of an effect from that of a cause. Parallel considerations can be made for the problem of deriving the belief state of a cause from that of an effect. In this case the following relations hold:

bel(P, true) = bel(Q, true) $*$ bell(Q $\xrightarrow[E]{}$ P)

bel(P, false) = bel(Q, false) $*$ bell(P $\xrightarrow[C]{}$ Q).

4.2 Global propagation

The global propagation procedure adopted in our Q/C-E networks aims at reproducing the following reasoning pattern: first, hypotheses explaining collected evidences are produced through evidential reasoning (along E-links only), then consequences of the drawn hypotheses are deduced through causal reasoning (along C-links only), and eventually an assessment of reasons for believing or disbelieving in each proposition is established. As a consequence *global propagation* is made up of the following 5 steps:

1. E-propagation: the initial belief states (acquired from the available evidences) are propagated along E-links through the entire Q/C-E network;
2. E-aggregation: at the end of step 1, any proposition generally has more than one associated belief state, derived through different E-propagation paths or from initial evidences; these belief states have therefore to be aggregated into a unique belief state representing how much the proposition is believed so far, on the basis of evidential reasoning only;
3. C-propagation: the belief states obtained in step 2 are propagated along C-links through the entire Q/C-E network;
4. C-aggregation: at the end of step 3, any proposition generally has more than one associated belief state, derived through different C-propagation paths; these belief states have therefore to be aggregated into a unique belief state representing how much the proposition is believed so far on the basis of the belief states propagated through causal reasoning only;
5. E-C aggregation: a final judgment about the belief state of the propositions is synthesized through aggregation of the results of E- and C-aggregation.

In propagating belief states throughout a C-E network one should pay particular attention to two points: cyclic dependencies and illegal inferences. The problem of *cyclic dependencies* [3] concerns the presence of loops composed by C-arcs or E-arcs within the network. It can be simply solved in our case by labelling each propagated belief state with the list of all previously visited nodes: propagation is then blocked when an already visited node is reached again. The problem of *illegal inferences*, as described in [10], concerns counter-intuitive conclusions that may be drawn if causal and evidential reasoning are indiscriminately applied. Focusing on the example proposed in [10], let us consider the three propositions: "It rained last night" (RAIN), "The sprinkler was on last night" (SPRINK), "The grass is wet" (WET). In an intuitive world model, both RAIN and SPRINK may cause WET and, as a consequence, RAIN and SPRINK may be evidentially hypothesized starting from WET. So, if both causal and evidential propagation is indistinctly applied one could erroneously deduce RAIN from SPRINK, which apparently contradicts common-sense. The problem can be solved if we assume that a belief state derived, for a node X, through a causal link ending in X can be propagated only along causal links leaving X (not along evidential links). That is, intuitively, if one knows the cause of a fact, he is not allowed to infer anything about other possible causes of the same fact (from causes one can only infer consequences). Illegal inferences are simply and naturally avoided in our approach since E- and C-propagation are carried out separately and can not be freely intermixed.

4.3 Aggregation

In order to formally define aggregation, let us introduce some preliminary concepts. We call *belief plane* the Cartesian plane of all possible belief states of a generic proposition P: the vertical axis represents the possible values of bel(P, true) and the horizontal axis those of bel(P, false). Values of both axes are represented by the nine belief labels E_1, .. E_9. Any belief state [bel(P, <u>true</u>) bel(P, <u>false</u>)] is represented by a point in the belief plane. Four points in the belief plane are worth special attention, since they convey a definite intuitive semantics; namely:

- the D point ([E_1, E_9]) which means totally disbelieved;
- the B point ([E_9, E_1]) which means totally believed;
- the U point ([E_1, E_1]) which means unknown;
- the C point ([E_9, E_9]) which means contradictory.

The above concept are shown in Figure 2. The C-U diagonal, called the *null diagonal*, is the locus of all belief pairs for which bel(P, true) = bel(P, false), so that one can not make any decision about the truth or the falsity of P. The greater the distance in the belief plane between the representation of a belief state and the null diagonal, the more it is possible to make a clear decision about a proposition. The points D and B, placed at the maximum distance from the null diagonal, are null-indecision points. Therefore, we can assign to each point of the belief plane a *decision strength*, proportional to its distance from the null diagonal: the maximum strength (say, 1) is assigned to both D and B, whereas the minimum - actually null - strength (say, 0) is assigned to the points on the null diagonal. Another remark is now appropriate. Let us consider, just as a starting point, the specific situation where to a proposition P two belief states are associated, namely: [E_1, E_1] (the U point) and [E_9, E_9] (the C point). Although these two belief states are equivalent with respect to the decision about the truth or the falsity of P (null decision strength), if one has to choose a single point in the belief plane representing the situation as a whole (i.e., the aggregation point), then he is inclined to choose the point C rather than U. In fact, the point C is more informed than U; it is grounded on a greater amount of evidence or, equivalently, on a more powerful belief. Therefore, in general, it can be argued that the total amount of belief carried by a belief state (i.e., belief in the truth of P plus belief in the falsity of P) has a role in choosing the aggregation point in the belief plane. Thus, we can assign to each point of the belief plane a *decision power*, proportional to its distance from the U point: the maximum power (say, 1) is assigned to C whereas the minimum power (say, 0) is assigned to U.

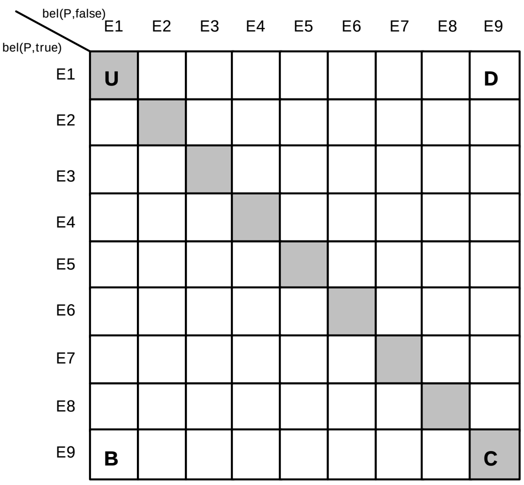

Figure 2: The belief plane.

Let us turn now to the main issue of aggregation.

For aggregation, two different approaches have been taken into account, namely: monotonic aggregation, and non-monotonic aggregation. Monotonic aggregation is used in step 4 (C-aggregation) while non-monotonic aggregation is used in the steps 2 and 5 (E- and EC-aggregation).

Monotonic aggregation, means that a belief state expressing disbelief must not produce a weakening effect when aggregated with a belief state expressing belief. This is according to common sense in C-aggregation. In fact, the disbelief in one of the causes of a certain effect, should not weaken the belief in the effect due to the presence of other causes. Therefore, aggregation is said to be monotonic, if a belief state resulting from a subsequent aggregation, can not have a distance from the B point greater than the distance of the belief state resulting from a previous aggregation. According to the intuitive concept introduced above, monotonic aggregation of a set of belief states associated to a proposition P (derived through different propagation paths), can be simply obtained by selecting in the set the belief state which is closest to the B point.

Non-monotonic aggregation, instead, does not imply any distance constraint related to results of subsequent aggregations. This can be achieved by taking into account both the distribution of different belief states and their number. For example, if there are two belief states associated to a proposition P which are in conflict, aggregating them requires being able to reach some sort of balance between them. Moreover, having just one belief state indicating that the proposition P is highly disbelieved is not the same, for a correct common sense judgment, as having four or seven of such belief states, assuming that they originate from independent evidence sources. This is according to common sense in E- and E-C aggregation. In fact, a belief based on a collected evidence, may be retracted if other contrasting evidences arise.

According to the intuitive concept introduced above, non-monotonic aggregation of a set of belief states associated to a proposition P (derived through different propagation paths), can be obtained by taking into account both the decision strength and the decision power of all the elements in the set. This can be done by defining a non-monotonic aggregation operator based on a principle of elastic equilibrium. More precisely, we first represent all the belief states of P in the correct positions on the belief plane and consider them as fixed points. Then, we provide ideal elastic springs connecting each of these points with a mobile point representing the (desired) result of non-monotonic aggregation (i.e., the aggregated belief state of P), and we define appropriate elastic constants (K_i) for such springs. For a generic point X in the belief, let us define K as:

$$K = \text{decision-strength}(X) + \text{decision-power}(X).$$

Finally, we compute the position of elastic equilibrium of P in the belief plane and interpret it as the aggregated belief state of P. More formally, given a set = {[t1 f1], [t2 f2], ..., [tn fn]}, the non-monotonic aggregation operator computes the belief state [tx fx], where tx and fx are the roots of the following equations:

$$K1*(tx - t1) + K2*(tx - t2) + ... + Kn*(tx - tn) = 0$$
$$K1*(fx - f1) + K2*(fx - f2) + ... + Kn*(fx - fn) = 0.$$

5 The ASTRA system

ASTRA[2] has been developed using KAPPA, a multiparadigm development tool running under Windows 3.1, on a 486 PC. In order to face the inherent complexity of the method human experts use to assess the state of a transformer and diagnose incipient faults, ASTRA has been designed as a distributed knowledge-based system, where separate agents, called *specialists*, cooperate to carry out the global

diagnostic task. Each specialist has its own knowledge base and is able to autonomously reason on (a part of the) input data and to draw partial conclusions. Specialist are divided in six competence areas, namely: user interaction, global problem, empirical, historical, structural, and causal.
Q/C-E networks have been used as the knowledge representation formalism of the specialists belonging to causal area. The formalism has been implemented using KAPPA's object-oriented features: node and links of the network are objects, respectively instances of the classes *Nodes* and *Links*. The class *Nodes* includes methods implementing the aggregation mechanisms, the class *Links* methods implementing local propagation. Global propagation is carried out through message passing among the objects representing the network.
All specialists have been developed and tested separately before being integrated within ASTRA. During development and testing, propagation and aggregation mechanisms have been repeatedly modified and refined. The version presented in this paper showed a good cognitive plausibility and matching with expert's reasoning features. The whole ASTRA system has been extensively tested with both simulated and real cases, showing a very good matching between expert's and system's judgments.

6 Conclusions

Q/C-E networks have been introduced in this paper as a formalism for reasoning under uncertainty in causal-evidential networks of propositions. Coherent methods for uncertainty propagation and aggregation in Q/C-E networks have been proposed. Q/C-E networks have been successfully experimented in the development of ASTRA, a knowledge-based system for preventive diagnosis of power transformers [2]. Future work on Q/C-E networks will deal with the following main issues:

- the extension of the concept of belief level in order to distinguish between uncertainty about the applicability and about the validity of a relation;
- the proposal and experimentation of new aggregation operators, based on different principles than elastic equilibrium;
- the explicit representation of the process of evidence-elicitation, considered as an integral part of uncertain reasoning;
- the investigation of uncertainty-based heuristics to be used for propagation in large networks, where exhaustive search can not be carried out.

Acknowledgements. The ASTRA project has been supported by the Consiglio Nazionale delle Ricerche, Progetto Finalizzato Robotica, Obiettivo PANDE. It has been developed in the period 1992-94 by a team including the University of Brescia, CILEA, and CESI s.p.a.

References

1. S.K. Andersen, K.G. Olesen, F.V. Jensen, and F. Jensen, HUGIN - A shell for building bayesian universes for expert systems, *Proc. 11th International Joint Conference on Artificial Intelligence,* Detroit, MI, 1989, 1080-1085.
2. P. Baroni, F. Cremonesi, G. Guida, S. Mussi, and S. Yakov, ASTRA: A knowledge-based system for state assessment, preventive diagnosis, and intervention planning of power transformers, *Proc. ISAP 94 Int. Conf. on Intelligent Systems Applications to Power Systems*, Montpellier,F, 463-470.
3. A. Bonarini, E. Cappelletti, and A. Corrao, Network-based management of subjective judgments: A proposal accepting cyclic dependencies, *IEEE Trans. on Systems, Man, and Cybernetics* SMC-22(5), 1992, 1008-1023.
4. P. Bonissone, Summarizing and propagating uncertain information with triangular norms, *Int. Journal of Approximate Reasoning* 1, 1987, 71-101.
5. P.P. Bonissone and K.S. Decker, Selecting uncertainty calculi and granularity: an experiment in trading-of precision and complexity, in L. Kanal and J. Lemmer (Eds.), *Uncertainty in Artificial Intelligence,* Elsevier, Amsterdam, NL, 1986, 217-247.
6. D.A. Clark, Numerical and symbolic approaches to uncertainty management in AI, *Artificial Intelligence Review* 4, 1990, 109-146.
7. D. Driankov, Uncertainty calculus with verbally defined belief-intervals, *International Journal of Intelligent Systems* 1, 1986, 219-246.
8. J. Pearl, Fusion, propagation and structuring in belief networks, *Artificial Intelligence* 29, 1986, 241-288.
9. J. Pearl, Bayesian decision methods, in G. Shafer and J. Pearl (Eds), *Readings in Uncertain Reasoning*, Morgan Kaufmann, San Mateo, CA, 1990, 345-352.
10. J. Pearl, Embracing causality in default reasoning, *Artificial Intelligence* 35, 1988, 259-271.
11. J.R. Quinlan, Inferno: A cautious approach to uncertain inference, *The Computer Journal* 26, 1983, 255-269.
12. G. Shafer, *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, NJ, 1976.
13. L. Zadeh, Fuzzy sets as a basis for a theory of possibility, *Fuzzy Sets and Systems* (1)1, 1978, 3-28.