

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

CONF-870832--3

DE87 014410

INCORPORATING "FUZZY" DATA AND LOGICAL RELATIONS IN THE DESIGN OF EXPERT SYSTEMS FOR NUCLEAR REACTORS

Michael A. S. Guth*

Instrumentation and Controls Division
Oak Ridge National Laboratory†
Oak Ridge, Tennessee 37831-6008

The submitted manuscript has been
authorized by a contractor of the U.S.
Government under contract No. DE-
AC05-84OR21400. Accordingly, the U.S.
Government retains a nonexclusive,
royalty-free license to publish or reproduce
the published form of this contribution, or
allow others to do so, for U.S. Government
purposes.

INTRODUCTION

This paper applies the method of assigning probability in Dempster-Shafer Theory (DST) to the components of rule-based expert systems used in the control of nuclear reactors. Probabilities are assigned to premises, consequences, and rules themselves. This paper considers how uncertainty can propagate through a system of Boolean equations, such as fault trees or expert systems. The probability masses assigned to primary initiating events in the expert system can be derived from observing a nuclear reactor in operation or based on engineering knowledge of the reactor parts. Use of DST mass assignments offers greater flexibility to the construction of expert systems in two important respects.

First, DST mass assignments have the advantage over classical probability methods of accommodating when necessary uncommitted probability assignments. Thus the DST probability framework can incorporate expert system inputs from imprecise or "fuzzy" data. Second, DST applied to the Boolean rules themselves leads to a probabilistic logic, where a given rule may be valid with probability less than unity: "fuzzy" logical rules.

When knowledge bases contain rules which may not always hold, or rules that occasionally must be operated upon with imprecise information, the DST mass assignments are shown to be a rigorous methodology for calculating probability assignments throughout the system.

The method of distributing probability to subsets of possible values of a random variable, rather than to individual sample points, was introduced by Dempster (1967), and was subsequently extended to the theory of belief functions by Shafer (1976, 1982, 1985, 1986).

*Subcontractor, this research was sponsored by appointment to the U.S. Department of Energy Laboratory Cooperative Postgraduate Research Training Program administered by Oak Ridge Associated Universities. Special thanks go to Bonnie Brummitt and the Publication and Information Processing Center staff for preparing this text.

†Operated by Martin Marietta Energy Systems, Inc., under contract DE-AC05-84OR21400 with the U.S. Department of Energy.

MASTER

Since 1985, DST has grown in popularity among AI researchers studying knowledge representation and reasoning under uncertainty, so that it now rivals Fuzzy Set Theory (FST) as a major uncertainty paradigm [Prade 1985]. The article by Lee, Grize, and Dehnad (1987) provides an excellent comparison of the theoretical and computational requirements of DST, FST, and classical probability theory.

Surprisingly, little work has been completed to date extending DST to truth-valued logic. Using DST, we will present a framework in which an analyst can represent the fact that he cannot determine whether some event has occurred or not. The event can be either true or false, but not both.

In classical two-valued logic, a statement or proposition is either true or false; similarly an event either occurred or did not occur. In three-valued logic, there exists a separate term denoting "unknown" or uncertain knowledge about the actual event.

Systems of Boolean equations form both fault trees and rule-based expert systems. This paper illustrates a simple set of Boolean equations depicting cause-consequence relations for a nuclear reactor or nuclear power plant to show how uncertainty over initiating events, subsequent results, or rules themselves can affect inferences with such models.

BACKGROUND FOR DST

Define the probability space (Ω, \mathcal{F}, P) , where Ω is a nonempty set of possible outcomes, \mathcal{F} is a σ -field of Ω , and P is a probability measure defined on \mathcal{F} . An element of Ω is called a sample point, while a subset of Ω is called an event. In DST probability is assigned to events, while in classical probability theory, probability is distributed among the sample points. Since a sample point is also a subset of Ω with one element, it can be treated as an event and have probability assigned to it in DST. This is how DST subsumes classical probability theory and generalizes it.

Mass Distribution Function

Consider the state space, Ω , with elements $\omega_i \in \Omega$. In classical probability theory, probability can be assigned to individual elements, ω_i , such that $\sum_i P(\omega_i) = 1$. The innovation of DST was to look to the power set 2^Ω , defined as the set of all possible subsets of Ω , and assign probability mass, $m_\Omega(\omega)$, to exclusive-OR (XOR) transformations of all $\omega \subset \Omega$. The mass distribution function of DST is a Borel measure satisfying the following two conditions: $0 \leq m_\Omega(\omega) \leq 1$ and $\sum \omega m_\Omega(\omega) = 1$. In the conventional DST notation ω_i is an element of Ω , while ω is both a subset of Ω and thereby an element of 2^Ω .

A third condition for $m_\Omega(\omega)$ is unique to DST. Implicit in the DST mass summation expression $\sum m_\Omega(\omega) = 1$ is the term $m(\phi)$, which denotes mass assigned to the null set. In DST, $m(\phi)$ is set equal to zero as a premise.

Three-Valued Logic

This paper considers a state space of $\Omega = \{\text{True}, \text{False}\}$ to evaluate events and the validity of propositions. The conventional selection of a state space in DST, also called the frame of discernment, has been chosen as the set of propositions (about a random variable's outcome) not the set of truth values. Therefore, this paper extends DST to a new domain, which we believe can bridge the gap between artificial intelligence and reliability analysis.

For our state space the power set $2^{\Omega} = \{\phi, T, F, (T,F)\}$ contains a term (T,F) denoting an exclusive-OR for an observation that could be either true or false, but not both. The (T,F) term comes from the DST uncommitted belief term in the power set. After renormalizing for ϕ , the effective power set becomes $2^{\Omega} = \{T, F, (T,F)\}$.

Modus Ponens

Definition 1: Material implication between A and B is expressed "If A, then B" and written " $A \rightarrow B$." Definition 2: The inverse of the logical relation " $A \rightarrow B$ " is " $A \rightarrow \neg B$." Definition 3: Modus ponens of the logical relation " $A \rightarrow B$ " is " $A \wedge (A \rightarrow B) \rightarrow B$." This paper determines the impact on modus ponens when a researcher receives information, that can be quantified into a probability, on the rule $(A \rightarrow B)$ itself. In binary logic, if a premise occurs, then the consequence is known as well. In designing expert systems, AI researchers use this logic and refer to "sure-fire" rules. If the premise occurs with some probability less than one, then the consequences would be expected to occur with the same probability. But what can be said if the rule itself is not strictly valid with probability one? The answer to this question can be found by first comparing the truth values emerging for the consequence, B, for given values of A and the rule itself $(A \rightarrow B)$.

Due to space limitations, we cannot provide here a formal justification of Table 1 based on DST. However, a formal derivation can be found in Guth (1987). Truth tables can also be constructed for the Boolean OR and AND operators that incorporate the three-valued logic. Again to conserve space, we will utilize the logic from these tables published elsewhere in the literature, e.g., Zimmerman (1985, p. 137), without reprinting them here.

Mass Assignment Equations

Denote the mass distribution for A as $m(A) = (a, b, c)$, where "a" denotes the mass assigned to the True element, b to the False element, and c to the XOR(T,F) element. Therefore, $a + b + c = 1$. Similarly denote the mass assignment for B by $m(B) = (x, y, z)$. Then we obtain the mass assignments for the Boolean AND and OR operations as follows:

$$m(A \wedge B) = (ax, b + y - by, az + cz + cx) , \quad (1)$$

$$m(A \vee B) = (a + x - ax, by, bz + cy + cz) . \quad (2)$$

From Table 1 if we let the mass on the rule $(A \rightarrow B)$ denoted by R and $m(R) = (r_1, r_2, r_3)$ then

$$m(B) = (ar_1, ar_2, 1 - ar_1 - ar_2) . \quad (3)$$

Table 1. Modus ponens truth table

		$A \rightarrow B$		
		B	T	F
		T	T	F
A	T	(T,F)	(T,F)	(T,F)
	F	(T,F)	(T,F)	(T,F)

NUCLEAR REACTOR EXPERT SYSTEM

4

To see the relevance of these two forms of uncertainty - uncommitted mass and probabilistic rules - on expert systems for nuclear reactors, consider the following example. Our objective is to conclude the probability that the top event will occur, and determine how uncertainty propagates through the model. Let a hypothetical expert system for a pressurized water reactor be summarized by the following six Boolean equations:

$$\text{Pressure Balanced} = \text{Power Source Normal} \wedge \text{Power Sink Normal} \quad (4)$$

$$\text{Power Sink Normal} = \text{Sec. Flow Normal} \wedge \text{Sec. Temp. Normal} \quad (5)$$

$$\text{Sec. Temp. Normal} = \text{Sec. Pumps Operating} \wedge \text{Condenser Vacuum Normal} \quad (6)$$

$$\text{Power Source Normal} = \text{Temp. Pri. Normal} \wedge \text{Pri. Flow Normal} \quad (7)$$

$$\text{Pri. Flow Normal} = \text{Pri. Pump Operating} \vee \text{Backup Pump On} \quad (8)$$

$$\text{Pri. Pump Operating} = \text{No Mechanical Failure} \wedge \text{Power Normal} \quad (9)$$

where \wedge denotes the Boolean AND operator and \vee denotes the Boolean OR operator, the abbreviation "Pri." refers to Primary, "Sec." to Secondary, and "Temp." to Temperature.

Each of the six Boolean equations can be interpreted as an "If, then" rule of a rule-based expert system. For example, Eq. (6) states "if the secondary pumps are operating AND the condenser vacuum is normal, then the secondary temperature is normal." Equation (8) states "If the primary pump is operating OR the backup pump is on then primary flow will be normal."

The equal sign in Eqs. (4) through (9) implies logical equivalence and could be replaced by the symbol \iff . Logical equivalence implies that the material implication (\rightarrow) works in both directions. Thus, e.g., Eq. (6) also states "If the secondary temperature is normal, then the secondary pumps are operating AND the condenser vacuum is normal."

To facilitate review of these equations, it will be convenient to introduce some short-hand notation for the events. Let the initiating events be denoted by letters A through F with intermediate events I_1 through I_5 leading to the top event of interest: Pres. Therefore, let Eqs. (4) through (9) be represented by

$$\text{Pres.} = I_1 \wedge I_3 \quad (4')$$

$$I_1 = A \wedge I_2 \quad (5')$$

$$I_2 = B \wedge C \quad (6')$$

$$I_3 = D \wedge I_4 \quad (7')$$

$$I_4 = E \vee I_5 \quad (8')$$

$$I_5 = F \wedge G \quad (9')$$

Thus the top event, Pres. represents "Pressure Balanced," and D represents "Temp. Pri. Normal," etc.

Let R_i denote rule i for $i = 4, 5, 6, 7, 8, 9$, and let the individual R_i correspond to Eqs. (4') through (9'), respectively. We will assign mass to the rules in the same way that mass is assigned to the initiating events. The R_i will be used as parameters for a study of modus ponens. The expert systems rules Eqs. (4') through (9'), coupled with individual R_i for the validity of each rule, are graphically depicted in Fig. 1.

The fact we have written Eqs. (4') through (9') with logical equivalence (\equiv) rather than material implication (\rightarrow) has nontrivial consequences for analysis with the tree in Figure 1. The equal sign implies that one can study effects propagating up or down the tree. With material implication, the logic would work in only the direction of the implication, e.g., only up or only down the tree.

Three different cases for the mass assignments on the initiating events and on the rules are shown in Table 2. For the first case the researcher can clearly observe whether the components A through G occur, but he faces some uncertainty over whether the rules are always valid. In Case II the researcher places uncommitted probability on both the components and the rules.

In Case III the researcher has precise information on both the components and the rules. Case III represents a conventional binary logic

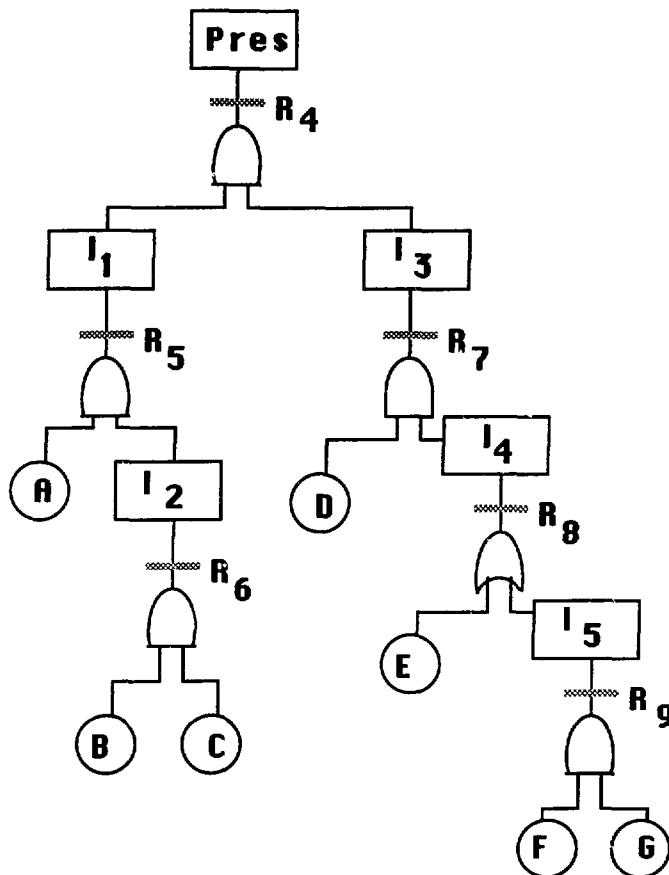


Fig. 1. Expert system rule tree.

Table 2. Mass assignments for alternative cases
(True, False, True-False)

Variable	Case I	Case II	Case III
A	(.9, .1, 0)	(.9, 0, .1)	(.9, .1, 0)
B	(.9, .1, 0)	(.9, 0, .1)	(.9, .1, 0)
C	(.7, .3, 0)	(.7, 0, .3)	(.7, .3, 0)
D	(.8, .2, 0)	(.8, 0, .2)	(.8, .2, 0)
E	(.7, .3, 0)	(.7, 0, .3)	(.7, .3, 0)
F	(.9, .1, 0)	(.9, 0, .1)	(.9, .1, 0)
G	(.8, .2, 0)	(.8, 0, .2)	(.8, .2, 0)
r ₄	(.8, 0, .2)	(.8, 0, .2)	(.8, .2, 0)
r ₅	(.9, 0, .1)	(.9, 0, .1)	(.9, .1, 0)
r ₆	(.6, 0, .4)	(.6, 0, .4)	(.6, .4, 0)
r ₇	(.8, 0, .2)	(.8, 0, .2)	(.8, .2, 0)
r ₈	(.7, 0, .3)	(.7, 0, .3)	(.7, .3, 0)
r ₉	(.8, 0, .2)	(.8, 0, .2)	(.8, .2, 0)

assignment, and one might expect the results from this case to eliminate mass assignments on the uncommitted (T,F) term. However, the application of modus ponens leads to uncommitted mass assignments on the consequence even in the binary mass assignment case.

The results for Cases I and II are identical and shown in Table 3. No mass is assigned to the false elements of the consequences in either of these two cases due to the nature of the modus ponens operator. Mass can only be assigned to the false element of the consequence if there exists positive mass on both the true element of the premise and on the false element of the rule. For Cases I and II no mass is assigned to the false element of the rule.

The results for Case III, given in Table 4, reveal the same mass assignment for the true element of the consequences; however, part of the mass previously assigned to the (T,F) element has been redistributed to the false elements. The probabilistic interpretations of initiating events shown in Case III have often been employed by the risk analysis community and form the basis for probabilistic risk assessment studies of nuclear reactor safety and exposure to radiation. Yet Case III also considers a probabilistic interpretation for rules, which thereby leads to the three-valued logic mass assignments for the consequences given in Table 4.

The mass assignments shown in Tables 3 and 4 were obtained using Eqs. (1), (2), and (3) with local computations for each intermediate event up the tree. See Shenoy and Shafer (1986) and Shafer, Shenoy, and

Table 3. Equilibrium mass assignments for Cases I and II

Event	True	False	True or false
Pres.	0.0958	0.0000	0.9042
I ₁	0.3062	0.0000	0.6938
I ₂	0.3780	0.0000	0.6220
I ₃	0.3910	0.0000	0.6090
I ₄	0.6110	0.0000	0.3890
I ₅	0.5760	0.0000	0.4240

Table 4. Equilibrium mass assignments for Case III

Event	True	False	True or false
Pres.	0.0958	0.0239	0.8803
I ₁	0.3062	0.0340	0.6598
I ₂	0.3780	0.2520	0.3700
I ₃	0.3910	0.0978	0.5112
I ₄	0.6110	0.2618	0.1272
I ₅	0.5760	0.1440	0.2800

Mellouli (1986) for a discussion on when local computation techniques can be employed.

For many rule-based expert systems the existence of possible nonunity mass assignments on initiating events or rules has been an important caveat preventing further inference. In DST mass assignments a rigorous methodology has been proposed to accommodate these probabilistic mass assignments and explain their impact on the mass assigned to the rules themselves.

The numerical results show that for some consequential events, e.g., I₄ and Pres., the majority of mass is assigned to (T,F). Moreover, assigning mass directly to the false element of the rule, as opposed to leaving "non-true" mass on the uncommitted belief term (T,F), does result in a redistribution of mass to the false element of the consequences.

For those rule-based systems which contain probabilistic or frequency interpretations of their components, DST mass assignments coupled with a probabilistic interpretation of rules represents a comprehensive framework for inference with imprecise information. Our experience with expert systems in robotics and nuclear reactor control confirms that many of the

expert systems extant use "decision rules" rather than strictly logical material implication rules. Thus we feel confident that the researcher interested in studying the effects of uncertainty in his expert system will find these methodologies useful.

DISCUSSION OF RELATED LITERATURE AND CONCLUSIONS

Dubois and Prade (1987) as well as Chatalic, Dubois, and Prade (1986) have written on DST extensions to expert systems. Both of these works employ the traditional set of propositions as a frame of discernment. Moreover, the latter paper sketches the expert system rules in the form of a "dependency tree" that resembles more of a semantic network than a fault tree with Boolean AND and OR gates.

A conference sponsored by Electric Power Research Institute (EPRI) in May 1987 had two sessions related to plant control: (1) Expert Systems Applications in Nuclear Power Plants and (2) Expert Systems Applications in Plant Diagnostics. One of the EPRI conference papers by Neuschaefer et al. (1987) reported development of a general expert system shell for plant diagnosis where the user supplies the knowledge base in the form of cause-consequence relations depicted in a fault tree. Their system proposes to account for uncertainty in instrument readings through conventional signal validation and filtering.

Skatteboe, Tangen, and Berge (1987) describe the use of alternative mathematical models in the construction of expert systems. These authors distinguish between "shallow" versus "deep" information, where shallow information corresponds to "rules-of-thumb" built into expert systems and deep information derives from actual statistical frequencies. The rules-of-thumb appear to be analogous to rules in our models that hold with a high probability, e.g., 0.9. Their same work also discusses qualitative reasoning based on imprecise quantitative measures.

Finally, DST has been incorporated into expert systems for structural damage assessment (Ishizuka, Fu, and Yao, 1981), and in a medical genetics diagnostic system (Gouvernet, Ayme, and Sanchez, 1982).

In conclusion, this paper has shown how uncertainty over the logical relations embodied in expert system rules can be mathematically represented and propagated through a system of Boolean equations. We have borrowed the concept of modus ponens to admit a probabilistic logic for the validity of rules, where some given rule may be valid with a probability less than one.

This paper has also demonstrated an application of the uncommitted belief term in DST to the possibility set {True, False}. The uncommitted term can represent component failures that may or may not have occurred, or the observer cannot tell whether it occurred. To the extent expert system researchers use "decision rules" and imprecise input data, the forms of uncertainty explored here should prove insightful to future modeling efforts.

REFERENCES

Chatalic, Philippe, Didier Dubois, and Henri Prade, (1986), "An Approach to Approximate Reasoning Based on Dempster's Rule of Combination," paper presented at RAI/IPAR '86 Conference on Robotics and Artificial Intelligence, June 18-20, 1986.

Dempster, Arthur P., (1967), "Upper and Lower Probabilities Induced by a Multi-Valued Mapping," Annals of Mathematical Statistics, Vol. 38, pp. 325-339.

Dubois, Didier, and Henri Prade, (1987), "On the Combination of Uncertain or Imprecise Information," L.S.I. Report #263, Universite Paul Sabatier, 118 route de Narbonne, 31062 Toulouse Cedex, France, submitted to Inter. Journal of Approximate Reasoning.

Gouvernet, J., S. Ayme, and E. Sanchez, (1982), "Approximate Reasoning in Medical Genetics," in Fuzzy Set and Possibility Theory: Recent Developments, R. R. Yager (Ed.), (New York: Pergamon Press).

Guth, Michael A. S., (1987), "Uncertainty Analysis of Rule-Based Expert Systems Using Dempster-Shafer Mass Assignments," Inter. Journal of Intelligent Systems, forthcoming.

Ishizuka, M., K. S. Fu, and J. T. P. Yao, (1981), "SPERIL-1: Computer-based Structural Damage Assessment Systems," CE-STR-81-36, School of Civil Engineering, Purdue University, West Lafayetta, IN, November 1981.

Lee, Newton S., Yves L. Grize, and Khosrow Dehnad, (1987), "Quantitative Models for Reasoning under Uncertainty in Knowledge-Based Expert Systems," Inter. Journal of Intelligent Systems, Vol. II, pp. 15-38.

Neuschaefer, Carl H., Peter W. Rzasa, Eugene Filshtein, Richard L. Burrington, and Robert Donais, "Application of C-E's Generic Diagnostic System to Power Plant Diagnostics," paper presented at EPRI conference on Expert Systems Applications in Power Plants, Boston, MA, May 1987.

Prade, H., (1985), "A Computational Approach to Approximate and Plausible Reasoning with Applications to Expert Systems," IEEE Trans. Pattern Anal. Mach. Intel., Vol. PAMI-7.

Shafer, Glenn, (1976), A Mathematical Theory of Evidence, (Princeton, N.J.: Princeton University Press).

(1982), "Bayes's Two Arguments for the Rule of Conditioning," Ann. Statistics, Vol. 10, pp. 1075-1089.

(1985), "Conditional Probability," International Statistical Review, Vol. 53, #3, pp. 261-277.

(1986), "The Combination of Evidence," Internat. J. Intelligent Systems, Vol. 1, No. 3, pp. 155-179.

(1986) with P. P. Shenoy, and K. Mellouli, "Propagating Belief Functions in Qualitative Markov Trees," Working Paper No. 186, School of Business, University of Kansas, Lawrence.

Shenoy, P. P. and G. Shafer, "Propagating Belief Functions Using Local Computations," IEEE Expert, Vol. 1, No. 3, pp. 43-52.

Skatteboe, Rolf, and Grethe Tangen, and Kaj Berge, (1987), "Models Applied in Knowledge Based Diagnosis," paper presented at EPRI conference on Expert Systems Applications in Power Plants, Boston, Mass.

Zimmerman, H.-J., (1985), Fuzzy Set Theory - and Its Applications, (Boston: Kluwer-Nijhoff Publishing Co. 1985).