An experimental comparison of knowledge engineering for expert systems and for decision analysis

Max Henrion,
Social and Decision Science,
and Engineering and Public Policy,
Carnegie Mellon University,
Pittsburgh, Pa 15213.

Daniel R. Cooley,
Department of Plant Pathology,
University of Massachusetts,
Amherst, Ma 01003.

Abstract

Decision analysis provides a set of techniques for structuring and encoding expert knowledge, comparable with knowledge engineering techniques for rule-based expert systems. In order to compare the expert systems and decision analysis approach, each was applied to the same task, namely the diagnosis and treatment of root disorders in apple trees. This experiment illustrates a variety of theoretical and practical differences between them, including the semantics of the network representations (inference net vs. influence diagram or Bayes' belief net), approaches to modelling uncertainty and preferences, the relative effort required, and their attitudes to human reasoning under uncertainty, as the ideal to be emulated or as unreliable and to be improved upon.[1]

I. Introduction

As schemes for representing uncertainty in AI proliferate and the debate about their various merits intensifies, [Kanal & Lemmer, 1986; Gale, 1986], it is becoming increasingly important to understand their relative advantages and drawbacks. One major axis of contention has been between proponents of various heuristic, qualitative, and fuzzy logic schemes, who argue that these are more compatible with human mental representations and consequently more practical to build and explain [Buchanan & Shortliffe, 1984; Cohen, 1985; Zadeh, 1986], and advocates of probabilistic schemes, who emphasize the virtues of being based on a normative theory of decision making under uncertainty [Pearl, 1985; Cheeseman, 1985; Spiegelhalter, 1986]. The latter have argued the advantages of approaches that are *coherent*, i.e. strictly consistent with the axioms of probability, over the earlier approximate Bayesian schemes developed for Mycin and Prospector [Duda *et al.*, 1976]. So far, comparisons have focused primarily on differences in theoretical assumptions [Bonissone, 1986; Horvitz, Heckerman & Langlotz, 1986; Henrion, 1987a], although there have been a few experimental studies which compare the performance of different uncertain inference schemes given knowledge formalized as a small rule-set [Tong & Shapiro, 1985; Wise & Henrion, 1986; Yadrick *et al*, 1986; Wise, 1986].

The informal experience of knowledge engineers and decision analysts alike suggests that choices about the structuring and encoding process that formalizes expert knowledge may have more impact on the final results than the numerical details of the uncertainty calculus employed. In the past, coherent probabilistic schemes have been criticized as intractable for significant practical applications, but recent developments have appear to have improved their practicality for construction and computation. These include *influence diagrams* [Howard & Matheson, 1984] and *Bayesian belief nets* [Pearl, 1986]. These are graphical tools which facilitate the qualitative structuring of uncertain knowledge and provide a framework for the numerical encoding of probabilistic relations in a form guaranteed to be coherent. The term *knowledge engineering* seems as appropriate for describing the activity of the decision analyst in building a probabilistic decision model to represent uncertain beliefs and preferences as it is for the construction of an expert system.

The purpose of this paper is twofold: First to illustrate the knowledge engineering process employing such a decision analytic approach, and second, to compare it with a rule-based expert system approach applied to the same problem. Since most readers will be more familiar with the latter approach, we shall provide greater detail on the former.

II. The experimental task

The task selected involved the diagnosis and treatment of root disorders of apple trees. We considered several causes of root damage, including water stress from waterlogged soil, cold stress from a severe winter, and the fungus, *phytophthora*. These problems are of major commercial significance to orchardists, and often lead to damage and destruction of apple trees. Moderate cases of phytophthora can be controlled by applying a fungicide. Other treatments include tiling and draining the area to control water damage, and bridge-grafting. If the damage has progressed too far, reducing apple production permanently by more than about 25%, the most efficient solution may be to destroy the trees and replant. The consultant plant pathologist uses a wide variety of evidence about the tree, environmental conditions, observable symptoms, and laboratory tests to diagnose

[1]This work was supported by the National Science Foundation under grant IST-8603493 to Carnegie Mellon.

the cause of root damage, and so recommend treatments. Figure 1 lists some of the elements of the problem.

Diagnoses: Phytophthora, cold stress, water stress.
Treatments: Fungicide, tiling and drainage, uproot and replant, wait and see.
Examples of evidential variables (with values):

- Winter cold episodes without snow cover (yes, no)
- Soil texture (light, moderate, heavy)
- Wetland vegetation (yes, no)
- Phytophthera resistant root stock (yes, no)
- Delineated root cankers (yes, no)
- Root tissue damage (None, little, moderate, severe)

Figure 1: Selected elements of the apple root problem

III. The Approaches

Two decision support systems were constructed to diagnose root disorders in apple trees and recommend treatment. One, which we will term the "ES model", was built by an knowledge engineer experienced in building knowledge-based expert systems. The other, which we will term the "DA model", was built by an experienced decision analyst. Both were developed on the basis of extensive interviews with a plant pathologist (D.R.C.), who has ten years experience as a specialist in this area. The initial structuring, encoding and implementation phases of the knowledge engineering process were carried out over an intensive four-day period, during which the two knowledge engineers alternated in working with the expert. The full implementation, testing, and refinement of the systems were completed over a longer time frame.

The ES model was implemented in KEE (Intellicorp) as a standard inference network with data-directed control. Diagnostic relationships are represented as rules giving the degree of belief in intermediate hypotheses and disorders based on Boolean combinations of data (evidential variables). Additional rules provide support for various treatments based on the diagnoses and other evidence.

The DA model employs an influence diagram to represent the expert's beliefs about how possible root disorders and treatments might affect the tree productivity and costs. It incorporates a Bayesian belief net to diagnose the disorders based on the available evidence. During the initial interview period, part of the influence diagram was used to construct a decision tree, which was implemented in Arborist (Texas Instruments) for preliminary analysis. Subsequently, the entire influence diagram including the diagnostic belief net was implemented using a combination of algorithms for propagating evidence through Bayes' nets [Pearl, 1986; Henrion, 1987b].

IV. Initial Structuring

The initial phase for both approaches was to identify the objects in the domain, that is, the root disorders, treatments, and evidential variables. Both knowledge engineers worked with the expert to draw directed graphs which represent qualitative evidential links between these elements. The first to be acquired was the influence diagram for the DA model. From this, and further discussion with the expert, an inference net was derived. These networks allow the decomposition of the expert's domain knowledge into separable local relationships. The initial influence diagram had 30 nodes, and the inference net had 25, of which 20 were common to both. Figure 2 shows a fragment of both networks superimposed for comparison.

Although they are topologically similar, there is a fundamental difference in the interpretation of the links. In the inference net, the direction of the links corresponds to the anticipated direction of inference, from evidence to disorders to treatments. In the influence diagram the direction of the links generally represents the believed direction of causal influence, for example abiotic stress increases susceptibility to phytophthora, and either of these can cause root tissue damage. Note that the influence diagram does not need to represent causal influences in full scientific detail (e.g. the physiology of how phytophthora produces root tissue damage) even when this is known, unless it seems likely to significantly improve the inference results. The influence diagram is also a way to express qualitative judgments about probabilistic independence: Two unlinked variables with a common cause (e.g. the phytopthora lab test and root tissue damage) are conditionally independent of each other and also of indirect antecedents (e.g. resistant root stock) given their immediate cause (e.g. phytophthora infection).

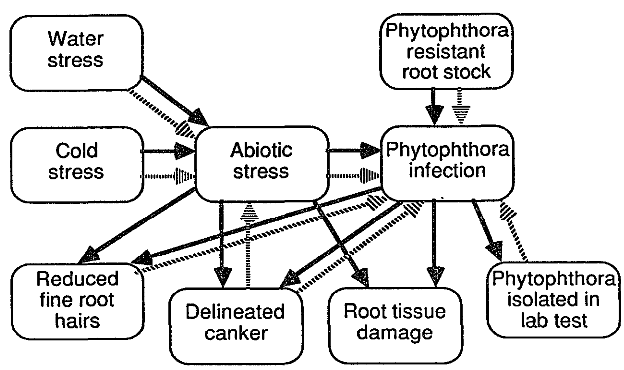

Figure 2: Fragment of influence diagram (solid arrows) with corresponding inference network (hatched arrows)

As Figure 2 shows, the links in the two representations may go in the same direction, such as where resistant root stock decreases susceptiblity to phytophthora infection *and* (therefore) provides evidence against phytophthora. More often they go in opposite directions, such as where observable symptoms are caused by a disease and therefore they are used as evidence for it. For example, phytophthora can cause delineated cankers, and so cankers are evidence for phytophthora.

For both approaches, the directed links help in the subsequent encoding of the relationships. For the ES approach, arrows converging on a node indicate a potential diagnostic rule with the antecedent nodes to appear in the condition and the destination node to appear in the action. For the DA approach, uncertain influences are encoded as conditional distributions with the antecedent nodes as the conditioning variables. Psychological research suggests it is generally easier to assess the probabilities of effects conditional on their causes (e.g. symptoms given diseases) than vice versa [Kahneman, Slovic & Tversky, 1982],

This ES model, like almost all rule-based expert systems, can only propagate evidence in the direction in which it is encoded, (no matter whether the inference is controlled by forward or backward chaining.) In contrast, in the DA model the influence diagram does not determine the direction of inference. By taking expectations over the conditions, this may be in the causal direction, or, by application of Bayes' rule, it may be the opposite, diagnostic direction, according to requirements of the application. For example, it is possible to determine the current probability of an unexamined symptom, based on observations of other symptoms, before deciding whether it is likely to be worthwhile to examine it.

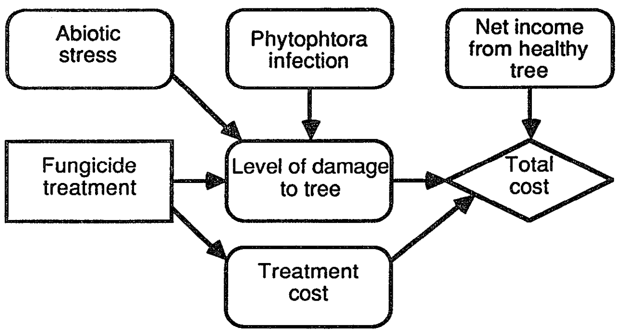

Figure 3: Influence diagram showing relation between root problems, treatment and outcome costs. The decision is enclosed in a rectangle, and the criterion variable is enclosed in a diamond.

V. Modelling costs and preferences

The DA approach developed an explicit quantitative model to estimate the costs and values of each combination of outcomes and treatments (Figure 3). Costs over multiple years were combined to obtain a discounted present value (a replacement tree takes 5 years to achieve full production). For the initial model, the orchardist was assumed to be risk neutral. The DA model computes the expected net cost of each treatment and recommends the treatment which minimizes it.

The ES model relies on heuristic rules for making inferences about what treatments to recommend based on the degrees of belief in the diagnoses without explicit consideration of costs. This approach seemed more natural for the expert, and was certainly much easier. Since the costs of treatment, tree replacement, and lost production for a given outcome do not vary greatly from one orchard to another, one can argue that such general rules may be widely applicable, somewhat analogous to the way it has been suggested that Mycin rules might be justified by decision analysis [Langlotz, Shortliffe, & Fagan, 1986]. However, as we shall see, at least in this case, formal analysis raises doubts about the adequacy of such informal analysis.

VI. Encoding relationships

In the ES model, the relationships in the inference network were encoded as sets of rules. The "degree of belief" in a hypothesis is one from the following ordered set of seven values: {confirmed, strongly-supported, supported, neutral, detracted, strongly-detracted, disconfirmed}. Each rule specifies the degree of belief in a conclusion based on combinations of its antecedents. For example, the rules in Figure 4 specify the degree of belief in phytophthora damage based on the values of six possible sources of evidence. These rules encode only those combinations of evidence thought by the expert to be important.

if phyto-resistance is low then supported
if reduced-fine-root-hairs is yes then supported
if reduced-fine-root-hairs is no then detracted
if cold-stress is at least supported or
    water-stress is at least supported then supported
if tissue-discoloration-below-soil is delineated-canker or
    tissue-discoloration-above-soil is delineated-canker
    then strongly-supported

Figure 4: Example diagnostic rules from the ES model for level of belief in phytophthora infection

Each uncertain influence is encoded as a probability distribution for the consequent, conditional on all its antecedents. To quantify each distribution, the expert was first asked for a verbal expression to get a rough idea, and then for explicit numbers. For example, evidence **v**, wetland vegetation, was judged to be causally influenced by hypothesis **w**, wet site. The expert judged it "quite probable" that wetland vegetation would be observed at a wet site, "impossible" if it was not a wet site. These judgments were then quantified as the two conditional probabilities, $p(v|\ w) = 0.7$, $p(v|{\sim}w) = 0$. Sensitivity analysis was used to examine the importance of accuracy in such assessments. In the vast majority of cases, a very rough assessment is perfectly adequate.

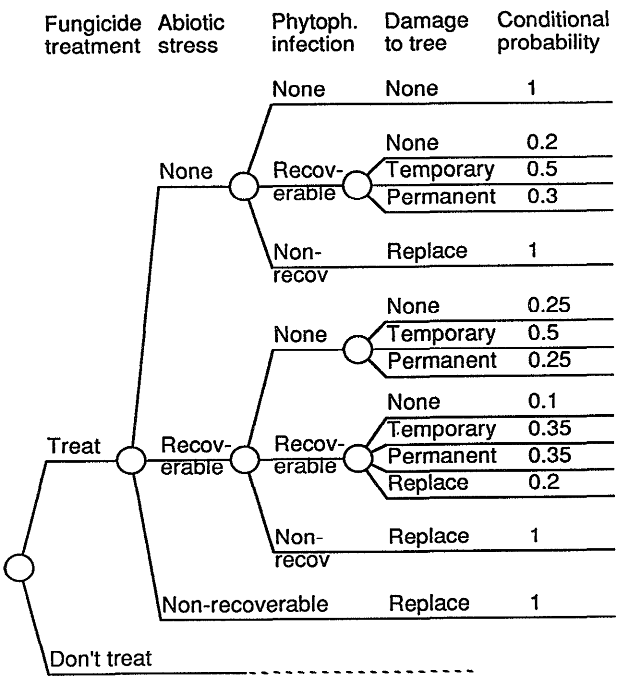

Figure 5: Partial decision tree for treatment decision

The number of parameters of the conditional distribution increases exponentially with the number of conditioning variables, but qualitative knowledge can usually reduce the assessment effort drastically. For example, the damage to the tree has 4 levels, conditioned on the severity of the phytophthora infection and abiotic stress, each at 3 levels, and the fungicide treatment decision (yes or no), giving a potential of 4 x 3 x 3 x 2 = 72 parameters to be judged. However, many are impossible, certain, or otherwise constrained by qualitative considerations, and there are actually only 12 different numbers requiring assessment. Figure 5 shows half of the decision tree whose terminal branches represent this conditional distribution.

VII. Testing and Refinement

In both approaches, initial implementations were tested to see if their conclusions were reasonable in the judgment of the expert, and the models were elaborated and tuned in the light of these tests. During the construction of the DA model, the use of conditional distributions to encode influences requires the expert to consider the impact of all possible combinations of evidence and decisions for each influence. The approach to encoding diagnostic rules for the ES model was much less demanding in its initial requirements of the expert. But it is consequently more likely to encounter combinations of events that had not been considered explicitly, and so requires more extensive testing and refinement.

The two approaches differ fundamentally in their response to unexpected results. For the ES model, rules were modified or added to obtain results that agreed more closely with the original expectations of the expert, since the primary goal was to emulate his judgment. In the DA approach, after initial rechecking of the relevant model structure and assessed probabilities, the probabilistic reasoning leading to the surprising conclusions was explained to the expert. If this lead him to accept them, the system was left unmodified. This was clearly illustrated by the following example.

Preliminary sensitivity analysis of the decision tree in Figure 5 showed that treating with fungicide had positive expected value, since the treatment is cheap (about $0.58 per tree) relative to the cost of replacing a tree (about $85). But the actual increment in expected value turned out to be very small, so small that it might often be outweighed by considerations not explicit in the model, such as environmental side-effects of the fungicide. This result was initially surprising to the expert, but examination of the model provided an explanation: The fungicide's effectiveness in controlling phytophthora was judged to be modest, and the probability that a tree is curable, i.e. that an infection is both present and not already beyond recovery is quite low. On reflection, he found this explanation convincing, and the result likely to be of considerable practical interest.

VIII. Discussion

No single knowledge engineer or decision analyst can claim that their approach is completely representative of all practitioners of their respective crafts. Certain aspects of both the approaches used here are somewhat atypical. The qualitative representation of uncertainty used in the ES model is less common than heuristic numerical schemes, such as Certainty Factors. The particular techniques applied here for evaluating influence diagrams are recent and not yet in general use. The size of both models and the effort devoted to their construction were quite modest. Nevertheless, several important points of comparison which are of general applicability to the two approaches, are clearly illustrated by this experiment.

In the initial structuring phase there are significant similarities between the approaches in the identification of the key elements, and the use of graphs to represent their interrelationships, but it is important to understand the fundamental differences in meaning between the inference network and influence diagram. Of course, research in expert systems has developed a rich array of techniques for knowledge represention and categorical reasoning that have not been a formal concern of decision analysis. It is specifically in the approaches to inference under uncertainty and decision making that the comparison is interesting. Here the main advantage of the ES approach is in the greater ease in initially encoding uncertain dependencies. This arises from the informal, heuristic nature of the language whether qualitative (as in this experiment) or quantitative, and the willingness to accept partial specifications, with inference rules conditioned on only the most salient combinations of evidence rather than the exhaustive combinations required for the DA model. The greater ease of encoding means that, for a given expenditure of effort, it is possible to build a system that deals with a larger number of sources of evidence, diagnoses and treatments than with the more rigorous DA approach.

The downside of this more relaxed approach is that the ES model is likely to require more extensive testing, debugging, and tuning to ensure that it performs adequately, and can handle common and important situations. Whether this additional effort will, in general, entirely cancel out its initial advantage is unclear from a single experiment and will depend on the criterion for adequacy.

Many in the AI community have believed coherent probabilistic approaches to be essentially intractable for problems involving the explicit representation of large bodies of expert knowledge. Part of this belief may stem from the traditional decision tree representation used by decision analysts, which grows exponentially with the number of uncertain events and decision variables. However, as illustrated, the influence diagram and Bayes' belief net provide tools for structuring and probabilistic inference which, if used judiciously, may have only linear complexity, albeit with a higher constant than the ES approach. A second common misgiving about the DA approach is the quantity of numerical judgments needed for assessing the probabilities. As we have seen, this may be greatly reduced by careful structuring and use of qualitative knowledge. Moreover, the vast majority of the numbers have small impact on the results, and so rough judgments will be adequate. Sensitivity analysis can help identify those few where significant assessment effort may be worthwhile.

Many of the advantages of the DA approach arise from its clearer separation of domain knowledge, obtained from the expert, and its general methods for inference under uncertainty based on Bayesian decision theory. The modelling of causal influences instead of inference rules provides an *isotropic* representation of domain knowledge with no preferred direction of inference [Henrion, 1987a]. Where the ES rules support inference only in the direction encoded, coherent Bayesian inference can perform causal, or diagnostic inference, as the occasion demands, operating on the same representation. The fact that causal models turn out to have advantages in representing uncertainty suggests an interesting relationship with other work on causal modelling for explanation and categorical diagnosis, for example in medical AI [Patil & Szolovits, 1981].

The treatment of discrepancies between the conclusions of the model and the expert illustrates a basic difference in philosophy. In the ES approach the performance of the expert is considered the "gold standard" which we seek to emulate. Discrepancies are therefore taken as a sign of a deficiency, and must be remedied by modifying or adding inference rules. The decision analyst also relies on the judgment of the expert, at least in those areas for which the expert has direct experience or knowledge. However, the decision analyst tends to be more impressed by the psychological findings on the limitations of human reasoning under uncertainty [Kahneman, Slovic & Tversky, 1982], and so is more skeptical about the expert's inferences beyond his immediate experience. If they disagree with the inferences made by the model, then the decision analyst may well prefer the results of the model. Indeed the possibility that the formal model may improve on the intuitive inferences of the expert is a major motivation for constructing it.

This was dramatically illustrated in the experiment by the low expected value of the fungicide treatment which the expert initially found counter-intuitive. After rechecking the assumptions and understanding the reasoning, he accepted its validity and modified his intuition. He found this an important insight likely to be of general value to other specialists in the area. The possibility of obtaining new results which go beyond current expert opinion requires a basis in some normative theory of decision making. Since the expert's original belief about the worth of the fungicide treatment was consistent with current opinion, and not liable to obvious empirical contradiction, such an insight would be hard to obtain by informal means.

IX. Conclusions

We have argued that the knowledge engineering effort required for a decision analytic approach is less than widely believed, and have demonstrated its feasibility for a significant application. However, although the influence diagram developed here is perhaps the largest yet reported, it remains two orders of magnitude smaller in scope than the largest expert systems reported (e.g. Internist/Caduceus), and it would be premature to make general claims about its capacity for major scale-up. Although decision analysis has been practiced successfully for over fifteen years, software to support influence diagram structuring and evaluation is still its infancy, though developing rapidly [Wiecha, 1986; Holtzman, 1985; Shachter, 1986; Henrion, 1987b]. The knowledge engineering effort for the decision analytic approach, will no doubt be significantly reduced, although is likely to remain somewhat greater than for rule-based expert systems.

However, there are considerable advantages to methods based on normative theory. It facilitates a consistent integration of causal and diagnostic inference. Uncertain beliefs and preferences are clearly differentiated. It provides a cleaner separation between domain knowledge and inference methods, and so may improve on the fallible reasoning of the expert. Whether these advantages outweigh the extra effort involved depends on the problem domain and task, and will remain partly a matter of taste. In the future, the dilemma may be resolved by systems that integrate ideas and techniques from both approaches to provide a richer range of options combining the advantages of each.

References

Bonissone, P. P. (1986). Plausible Reasoning: Coping with Uncertainty in Expert Systems. In *Encyclopedia of AI.* John Wiley & Sons.

Buchanan, B. G. & Shortliffe, E. H. (1984). *Rule-based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project.* Addison-Wesley, Reading, Mass.

Cheeseman, P. (1985). In Defense of Probability. *Proceedings of 9th International Joint Conference on AI.* Los Angeles, Ca.

Cohen, P.R. (1985). *Heuristic Reasoning about Uncertainty: An AI Approach.* Pitman: Boston.

Duda, R. O., Hart, P. E., and Nilsson, N. J. (1976). *Subjective Bayesian Methods for Rule-Based Inference Systems* (Tech. Rep. SRI Technical Note 124). SRI International.

Gale, W. A. (ed.). (1986). *AI and Statistics.* Addison-Wesley: Reading, MA.

Henrion, M. (1987). Uncertainty in Artificial Intelligence: Is probability epistemologically and heuristically adequate? In J. Mumpower (Ed.), *Expert Systems and Expert Judgment.* NATO ISI Series: Springer-Verlag.

Henrion, M. (1987). Propagation of Uncertainty by Logic Sampling in Bayes' Networks. In L.N. Kanal & J. Lemmer (Eds.), *Uncertainty in Artificial Intelligence.* Amsterdam: North-Holland.

Holtzman, S. (1985). *Intelligent Decision Systems.* Doctoral dissertation, Engineering-Economic Systems, Stanford University,

Horvitz, E.J., D.E. Heckerman & C.P. Langlotz. (1986). A framework for comparing alternative formalisms for plausible reasoning. *Proceedings of AAAI-86: Philadelphia, PA.* American Association for Artificial Intelligence.

Howard, R.A. & J. E. Matheson. (1984). Influence Diagrams. In Howard, R.A. & J. E. Matheson (Eds.), *The Principles and Applications of Decision Analysis: Vol II.* Strategic Decisions Group, Menlo Park, CA.

Kahneman, D.; Slovic, P. & Tversky, A. (1982). *Judgment under Uncertainty: Heuristics and Biases.* Cambridge: Cambridge University Press.

Kanal, L.N. & Lemmer, J. (eds.). (1986). *Uncertainty in Artificial Intelligence.* North-Holland: Amsterdam.

Langlotz, C.P., E.H. Shortliffe, & L.M. Fagan. (1986). Using decision theory to justify heuristics. *Proceedings of AAAI-86.* American Association for Artificial Intelligence.

Patil, R.S. & Szolovits, P. (1981). Causal understanding of patient illness in medical diagnosis. *Proceedings of IJCAI-81: Vancouver, Canada.* , International Joint Committe for AI.

Pearl, J. (1985). *How to do with probabilities what people say you can't* (Tech. Rep. CSD-850031). Cognitive Systems Laboratory, Computer Science Department, University of California, Los Angeles, CA.

Pearl, J. (1986). Fusion, Propagation, and Structuring in Belief Networks. *Artificial Intelligence, 29*(3), 241-288.

Shachter, R.D. (1986). David: Influence Diagram Processing System for the Macintosh. In *Proceedings of Second Workshop on Uncertainty in Artificial Intelligence.* AAAI.

Spiegelhalter, D. J. (1986). A statistical view of uncertainty in expert systems. In W. Gale (Eds.), *AI and Statistics.* Addison-Wesley: Reading, MA.

Tong, R. M., Shapiro, D. G. (1985). Experimental Investigations of Uncertainty in a Rule-based System for Information Retrieval. *International Journal of Man-Machine Studies* .

Wiecha, C. (1986). *An empirical study of how visual programming aids in comprehending quantitative policy models.* Doctoral dissertation, Carnegie Mellon University, Department of Engineering and Public Policy,

Wise, B.P. (1986). *An experimental comparison of Uncertain Inference Systems.* Doctoral dissertation, Carnegie Mellon University, The Robotics Institute and Department of Engineering and Public Policy.

Wise, B.P. & Henrion, M. (1986). A Framework for Comparing Uncertain Inference Systems to Probability. In Kanal, L.N. & Lemmer, J. (Eds.), *Uncertainty in Artificial Intelligence.* North-Holland: Amsterdam.

Yadrick, R. M., Vaughan, D. S., Perrin, B. M., Holden, P. D., Kempf, K. G. (1986). Evaluation of Uncertain Inference Models I: PROSPECTOR. In *Second Workshop on Uncertainty in AI.* AAAI.

Zadeh, L. (1986). Is probability theory sufficient for dealing with uncertainty in AI? A negative view. In L.N. Kanal & J. Lemmer (Eds.), *Uncertainty in Artificial Intelligence.* Amsterdam: North-Holland.