From: AAAI-86 Proceedings. Copyright ©1986, AAAI (www.aaai.org). All rights reserved.

Using Decision Theory to Justify Heuristics

Curtis P. Langlotz, Edward H. Shortliffe, Lawrence M. Fagan

Medical Computer Science Group
Knowledge Systems Laboratory
Medical School Office Building
Stanford University Medical Center
Stanford, California 94305

ABSTRACT

We present a method for using decision theory to evaluate the merit of individual *situation -> action* heuristics. The design of a decision-theoretic approach to the analysis of heuristics is illustrated in the context of a rule from the MYCIN system. Using calculations and plots generated by an automated decision making tool, decision-theoretic insights are shown that are of practical use to the knowledge engineer. The relevance of this approach to previous discussions of heuristics is discussed. We suggest that a synthesis of artificial intelligence and decision theory will enhance the ability of expert systems to provide justifications for their decisions, and may increase the problem solving domains in which expert systems can be used.

I INTRODUCTION

The rule-based expert system [1], [2] is an established and widely used artificial intelligence paradigm. The rules in such expert systèms are often described as *heuristics*, which encode the experiential knowledge of experts for use in decision support systems. The capability of some expert systems has been shown to be comparable to experts (see for example [3], [4], [5]). However, the task of building expert systems, or *knowledge engineering*, has yet to be characterized in terms that allow the assessment of the merits of individual heuristics.

Nevertheless, previous attempts to characterize heuristic rules have led to insights intended to help knowledge engineers craft heuristics that lead to high performance. For example, Clancey [6] asked the question: What kinds of arguments justify rules and what is their relation to a mechanistic model of the domain? By analyzing a rule used in MYCIN [7], he demonstrated that a heuristic can be broken into smaller and smaller inference steps that support it. Lenat [8] also investigated the nature of heuristics, and asked: What is the source of the power of heuristics? He hypothesized that heuristics derive some of their power from regularity and continuity in the world. To illustrate this point, he provided qualitative plots of the power or utility of a heuristic against characteristics of the task domain. Smith [9] describes an expert system that explicitly represents justifications for heuristic rules and uses those justifications to guide knowledge base refinement. However, this system makes its decisions about the causes of system errors based on rule type (e.g., definitional, theoretical, statistical, or default), not based on the measures of certainty associated with individual rules. Gaschnig [10] quantitatively assessed the performance of heuristics used in search, but only by observing repeated trial executions of a search program with different heuristics.

When a rule is placed in the knowledge base, often no formal analysis is made to ascertain the power of the rule and the magnitude of its effect on system performance. Until a blinded evaluation study is completed, the system builder must assume that the heuristic suggested by the domain expert is appropriate for most or all cases the system will encounter. But heuristics almost always represent significant tradeoffs between possible costs and benefits, and the appropriateness of a heuristic may therefore often be argued.

In order to make a reasoned decision about a heuristic that recommends an action, it is important to explicitly consider both the likelihood and the desirability of the consequences of the action. Consequently, we argue for an analysis of heuristics based on the synthesis of artificial intelligence and decision theory. Decision theory can be used to combine explicitly expressed probabilities (likelihoods) and utilities (desirabilities) to decide between competing plans of action. It is an axiomatized method for making decisions which recommends the course of action that *maximizes expected utility*. The expected utility of a given plan is expressed as follows:

$$\text{Expected Utility} = \Sigma_i \; p(O_i) \times U(O_i)$$

where $p(O_i)$ is the probability of the ith outcome of executing the plan, and $U(O_i)$ is the utility of the ith outcome. This concept has been promoted by Savage [11], who defends subjective probabilities to represent uncertainty, and a utility function to represent preferences. Raiffa [12] and Howard [13] both provide a thorough introduction to decision theory.

Decision theory has been suggested as an adjunct to planning systems. Jacobs [14] and Coles [15] described robot planning systems that used AI techniques to generate plans, coupled with decision-theoretic techniques to compare plans based on costs and risks associated with planning operators. Feldman [16] described a similar framework that was used to solve a more realistic version of the "monkey and bananas" problem. Slagle [17] describes an interactive planning system that uses a predictive model of military damages to rank competing plans for allocation of military resources. We have also described a medical problem that could not be solved without explicit quantification of the uncertainties and tradeoffs involved [18].

We believe that precise definitions of both the application area and the notion of heuristic power or utility can provide important information to the knowledge engineer. For example:

1. How often will a heuristic be incorrect?
2. How does system performance change when a heuristic is added?
3. What serves as appropriate support or justification for a heuristic?

In an attempt to answer these questions, we first define our notion of a heuristic. Then we show how information generated by a decision analysis tool developed on a Xerox

*Support for this work was provided by the National Library of Medicine under Grants LM-04136 and LM-04316, the National Science Foundation under Grant IST83-12148 and the Division of Research Resources under Grant RR-01636. Computing facilities were provided by the SUMEX-AIM resource under NIH Grant RR-00785, by the Xerox Corporation, and by Corning Medical. Dr. Shortliffe is a Henry J. Kaiser Family Foundation Faculty Scholar in General Internal Medicine.

1100-series LISP machine can be used to analyze a particular heuristic. Next, we show how our analysis relates to earlier analyses of heuristics. Finally, we discuss the implications of our analysis for expert systems.

II THE FORMAL ANALYSIS OF A HEURISTIC

We will adopt the definition of a heuristic proposed by Lenat [8]. He defines a heuristic as "a piece of knowledge capable of suggesting plausible actions to follow or implausible ones to avoid." We have chosen to concentrate our discussion on a frequently cited heuristic rule from the MYCIN system [7], shown in Fig. 1.

```
If: 1) The therapy under consideration is
       tetracycline
    2) The age (in years) of the patient is
       less than 8
Then: There is strongly suggestive evidence
      (.8) that tetracycline is not an
      appropriate therapy for use against
      the organism
```

Figure 1: The MYCIN tetracycline heuristic, slightly simplified for illustration purposes.

Fig. 2 shows one possible collection of *support knowledge* that Clancey proposed as justification for this heuristic. Let's analyze this chain of four support rules in more detail. The first three inferences indicate how each event influences the occurrence of the next. But the final inference suggests a decision for *action* that is based on the previous inferences. No matter how fine the granularity of the reasoning, one rule in the chain will always recommend action based on the situation. That rule represents a compiled plan for action that will have wide ranging consequences. For example, avoiding the administration of tetracycline has the advantage of essentially eliminating the possibility of stained teeth, but it has the disadvantage of creating the need for another drug which may have a weaker therapeutic effect and other undesirable side effects. In fact, a widely used physician's reference book [19] states that tetracycline should not be used in children under age 8 *unless other drugs are not likely to be effective or are contraindicated.* In other words, there is a tradeoff between the undesirability of possible staining and the desirability of increased effectiveness.

```
tetracycline in youngster
 => chelation of the drug in growing bones
     => teeth discoloration
         => undesirable body change
             => don't administer tetracycline
```

Figure 2: A justification for the tetracycline heuristic in MYCIN from [6].

While these tradeoffs are important when deciding whether or not to recommend tetracycline, they have relevance in other settings. For example, they are essential to an intuitive justification of the conclusion not to recommend tetracycline. A justification for such a decision might be: "Although tetracycline is more likely to cure this infection, that is outweighed by the fact that tetracycline is likely to cause substantial dental staining."

These tradeoffs are important when deciding whether or not to include the tetracycline rule in an expert system. Although the addition of a certainty factor (in this case 0.8) is designed to allow other heuristics to override the recommendations of this one, it does not explicitly represent the circumstances under which the heuristic should be invalidated. Because the tradeoffs are not represented explicitly, the rule cannot recognize the characteristics of an unusual decision situation and sometimes select tetracycline in spite of possible cosmetic problems, just as an expert would. For example consider the following cases, for which the value of the tetracycline heuristic might be questioned:

- What if the infecting organism were resistant to all drugs except tetracycline?
- What if the only undesirable bodily change that tetracycline caused was minor intestinal distress?
- What if the probability of staining due to tetracycline was only 1 in 100? 1 in 1000?

III PROBLEM FORMULATION

This section describes a decision-theoretic method for representing the tradeoffs that underly the rule in Fig. 1. When formulating a problem in decision-theoretic terms, three questions must be answered:

1. What alternative plans are available?
2. What might occur if each of those actions were carried out?
3. What is the utility of each possible outcome?

We will examine a specific case in which the tetracycline rule might apply, and show how the results of the analysis can be generalized for use in building expert systems. Our problem will be constrained by considering only two alternative plans, and only a few possible outcomes of those plans. Although not shown here, the process of finding a small number of candidate plans can be automated [20].

The case concerns an 8 year old male who has a urethral discharge (an indication of possible urethral infection) but in whom cultures have shown no evidence of bacterial infection. In such cases the urethritis may be caused by organisms that cannot be cultured easily (non-specific urethritis, or NSU) or it might be related to a non-infectious process causing urethral inflammation. In adults with such symptoms, it is common to treat with tetracycline since it is usually effective in NSU and can help assure relief from discomfort. In a child, however, the risk of tetracycline, as summarized above, cannot be totally ignored. The specific question that must be decided is: Should this young patient be treated with tetracycline, or with the second choice drug, erythromycin? Erythromycin, unlike tetracycline, has no significant side effects except occasional nausea, but has the disadvantage that it is slightly less likely to cure the NSU.

To formulate a decision-theoretic representation of the problem, first the available actions must be enumerated: in this case to administer tetracycline or to administer erythromycin. Then the consequences of each action must be explored. In this case, if either action is performed, there are two possible scenarios to consider: The patient either has NSU or has a non-infectious urethritis. If the urethritis is infectious, then the tetracycline will be more likely to cure the infection than erythromycin. If it is non-infectious, then the drugs will have no therapeutic effect (except for a small placebo effect that is the same for both drugs). Finally, the undesirable side effects of tetracycline must be considered. Regardless of the outcome of tetracycline therapy, there is a definite chance that dental staining will occur. In summary, there are four pertinent outcomes that should be considered in delineating treatment options: CURE/NO STAINING, NO CURE/STAINING, CURE/STAINING, NO CURE/NO STAINING.

Once the decision options and their possible consequences have been enumerated, decision analysts conventionally represent the problem as a decision tree*. In Fig. 3 we see the tree that represents the decision problem described above. Each path through the decision tree represents one possible combination of actions and consequences that might occur. For example, the top branch represents the following chain of events: The patient had an infectious urethritis, was given

*Although decision trees are still the predominant representation convention, some members of the decision analytic community are increasingly attracted to an alternative representation called *influence diagrams* [21]. The intuitive, modular, characteristics of influence diagrams are similar to the AI representation techniques from which they are derived [22].

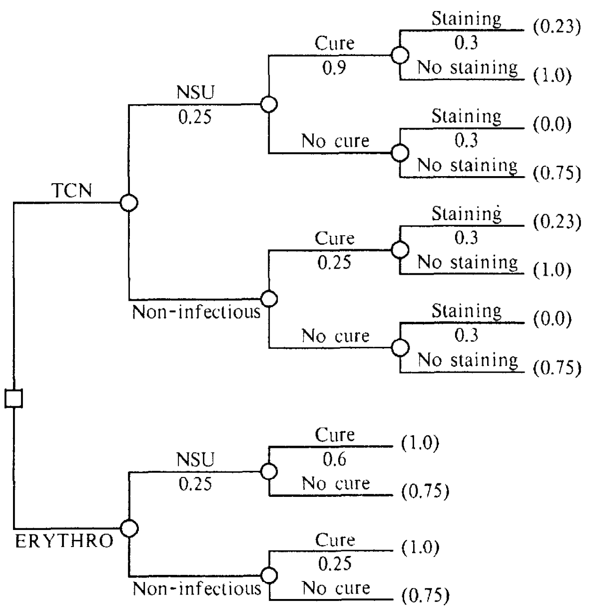

Figure 3: A decision tree that represents the decision between tetracycline and erythromycin for treatment of possible NSU. Square nodes are *decision nodes*. Branches emanating from decision nodes represent actions among which a choice must be made. The remaining nodes are *chance nodes*, whose branches represent all of the possible outcomes that might occur. The tree is labeled with the probabilities and utilities as assessed from a physician. TCN = tetracycline, ERYTHRO = erythromycin, NSU = non-specific urethritis.

tetracycline, which cured the disease, but dental staining resulted.

IV PROBABILITY AND UTILITY ASSESSMENT

For many tradeoffs, there is a point at which a small chance of a highly undesirable outcome will be equally preferred to a high likelihood of a mildly undesirable outcome. The point where this equivalence occurs may be dependent on precise expert assessments of probability and utility.

Although these assessments may be subject to some inaccuracies and biases [23], we will see in the next section that we need not utilize the precise values of these numbers to justify a decision. We need only show that large variations from the assessed value will not affect the decision. To assess the relevant probabilities, the following questions will be asked in the context of the particular patient:

1. What is the probability that tetracycline will cure non-specific urethritis (NSU)?
2. What is the probability that erythromycin will cure NSU?
3. What is the probability that dental staining will occur if tetracycline is administered to this patient?
4. What is the probability that this patient has an infectious NSU?
5. What is the probability that either drug will cure a non-infectious urethritis (through a placebo effect)?

To assess the utility of each of the four outcomes, explicit quantitative comparisons must be made among them*. The *standard gamble* is used to assess the utility of outcomes by converting a utility question into a probability question. Since utilities are relative quantities, it is conventional to assign the worst outcome a utility of 0.0 and the best outcome a utility of 1.0. Outcomes whose utilities are intermediate are assessed by asking the expert what gamble between a bad outcome and a good one would be equally preferable to the certainty of the intermediate outcome. The response to this question uniquely determines the relative desirability of the intermediate outcome. For example, if the expert were indifferent between guaranteed NO CURE/NO STAINING and a gamble with 1 chance in 4 of NO CURE/STAINING (utility = 0.0) and 3 chances in 4 of CURE/NO STAINING (utility = 1.0), then NO CURE/NO STAINING can be assigned a utility of 0.75. An analogous standard gamble question can be devised to find the utility of the CURE/STAINING outcome.

Fig. 3 shows the values of the parameters of the model as assessed from a physician. For this decision tree, the expected utility of administering tetracycline is 0.63, and the expected utility of administering erythromycin is 0.83**. Therefore, it would seem that in this case erythromycin is "better" than tetracycline, consistent with the original heuristic statement shown in Fig. 1. But how certain should we be of this conclusion? What does a difference of 0.2 utility units mean?

Since there is uncertainty about the values of the probability and utility parameters even when considering an individual patient, many object that probability assessments require of the expert a level certainty that cannot, in reality, be obtained. Additional uncertainty is introduced when generalizing to an entire set of cases to which an expert system will be exposed. To address these concerns, decision-theoretic techniques have been devised to answer the following question: If the value were different than the one provided by the expert, how likely would it be to affect the decision? The principal tool for this purpose, *sensitivity analysis*, is described in the next section.

V SENSITIVITY ANALYSIS

Identifying the variables to which a heuristic is sensitive can help determine the merit of the heuristic, can help provide an adequate justification for the heuristic, and can help direct ongoing knowledge acquisition efforts to those areas where further investigation is needed. To quantitatively assess the effect of changes in a variable, *one-way* sensitivity analysis is frequently employed. It determines how much one parameter in the decision model must vary before the optimal decision changes. Consider, for example, how the utility of administering each drug might change with changes in the probability of dental staining. A plot generated by such an analysis is shown in Fig. 4. The point at which the utilities of the plans are the same is called the *threshold* value. In this case, the threshold occurs when the probability of dental staining is equal to 0.025, quite a distance from the original assessed value of 0.3. If the threshold value were nearly equal to the assessed value, further analysis or data collection may be necessary to reach a decision.

The frequency with which a particular decision will be optimal depends in part on the chance that such a parameter will vary beyond the threshold. If the parameter was not known with great certainty, or if it varied considerably from

*There are two exceptions to this statements. First, in some simple problems such as the one we consider here, the dominance of one alternative can be proven solely from qualitative assertions about the relative utilities of the outcomes [24]. Second, since the number of outcomes that must be assessed grows rapidly with the size of the problem, not all these assessments are actually made in problems more complex than the one we consider here. Instead, decision analysts look for independent measures of utility that can be combined in an additive *utility model*, and make assessment of the parameters of that model.

**Note that these expected utility values are *not* certainty factors.

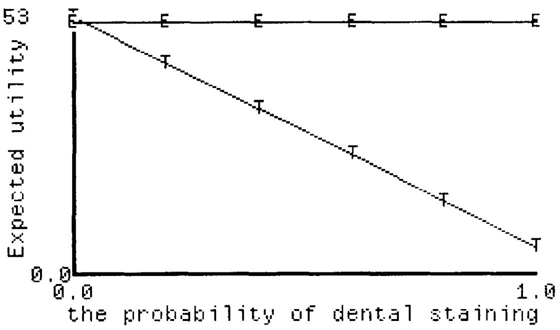

Figure 4: The results of a one-way sensitivity analysis of the tree shown in Fig. 3. The expected utility of each decision option (on the vertical axis) is plotted against the likelihood of dental staining due to tetracycline. T = tetracycline, E = erythromycin.

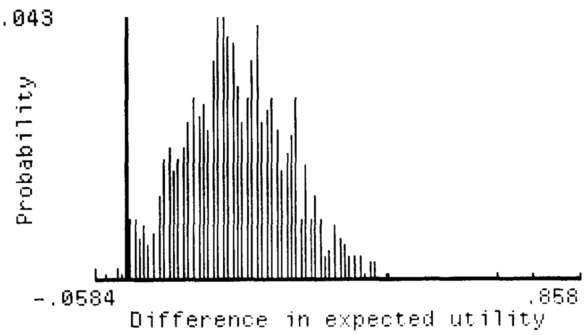

Figure 5: The results of a Monte Carlo simulation of the difference in expected utility between erythromycin and tetracycline (horizontal axis). The frequency with which each expected utility value occurred is plotted on the vertical axis.

case to case, the second choice might be the optimal choice in a substantial minority of instances. It is for this reason that decision analysts assess the approximate probability distributions of all sensitive parameters. For example, the expert can be asked to specify a range in which the value can be expected to fall half the time. This specifies an approximate distribution for the parameter [25]. This distribution represents the state of the expert's knowledge about the parameter.

Once such an assessment has been made, it is straightforward to find the probability that the value will fall beyond the threshold (by integrating the distribution up to that point). According to the assessed distribution of the prôbability of dental staining*, a value beyond the threshold occurs in less than 1 in 100,000 patients.

However, this value is only a lower bound on the probability of error, since a one-way sensitivity analysis assumes that only one variable at a time deviates from its mean value. It is possible that interactions between variables could cause substantially greater errors that would remain undetected by one-way sensitivity analyses. To address these concerns, more comprehensive sensitivity analyses have been developed, such as multi-way and Monte Carlo sensitivity analyses [26]. The Monte Carlo sensitivity analysis, in particular, provides an important metric for the evaluation of a heuristic. In Monte Carlo analysis, a value is randomly selected from the distribution of each relevant parameter, and the expected utility of the decision is computed for that random set of parameter values. This process is repeated many times to obtain an estimate of the distribution of the result. Fig. 5 shows the results of a Monte Carlo simulation of the difference between two competing alternatives. From this distribution, a number of useful quantities can be obtained. Since the figure shows the distribution of the *difference* between erythromycin and tetracycline, any negative value represents a set of parameter values for which tetracycline would be optimal (in direct contradiction to the original heuristic). The proportion of negative values represents the error rate of the heuristic and can serve as a useful indicator of the *power* of the heuristic.

*The distribution is not shown here. The knowledge was represented by a β-distribution with parameters $R = 6$ and $N = 20$. There are important theoretical reasons for selecting β-distributions, but these will not be presented here.

VI IMPLICATIONS FOR EXPERT SYSTEMS

Although assessing the quantities for a decision-theoretic analysis requires extra effort (in this case, seven quantities must be assessed), the required effort yields substantial advantages. For example, both the knowledge engineer and the domain expert are forced to be explicit about the population of cases for which the system is designed. This allows the identification of those cases to which the heuristic may not be useful, and the quantification of expected change in system performance.

We have shown that it might indeed be appropriate in some cases to administer tetracycline to a young child. However, Clancey's analysis leaves to intuition the notion that the undesirable bodily changes caused by tetracycline are sufficiently severe to outweigh the increased effectiveness of tetracycline. He makes explicit the causal chain of reasoning that indicates an undesirable bodily change may take place, but does not explicitly represent the tradeoffs between that undesirable change and the possibility of the poor consequences of not being treated by tetracycline.

As we discussed in section II, there are several possible scenarios in which the chain of rules in Fig. 2 might not justify the heuristic. Why, then, was MYCIN so successful? For a case similar to the ones addressed by MYCIN, the results of the Monte Carlo analysis indicate it is highly unlikely that a given patient would be better off with tetracycline. Furthermore, MYCIN was evaluated by comparing it to experts who also may choose to use the tetracycline heuristic for decision making, even though it does not always lead to the optimal decision. In any case, since other drugs are often as effective, any possible error would not be serious. These features may not be present in other less forgiving problem solving settings (e.g., aminoglycoside antibiotics are frequently used to treat for an infection with the organism *pseudomonas*, despite a high chance of nephrotoxicity).

VII CONCLUSION

We have demonstrated a decision-theoretic approach to the analysis of heuristics. The informational needs of this analysis technique can be provided through a process similar to conventional knowledge engineering. The concise fashion in which the problem is stated, together with the extra information obtained in the knowledge acquisition process, supplies tools for analyzing the performance of an individual heuristic. This decision-theoretic approach may help to augment the capabilities of expert systems.

We recognize that for complex problems, a decision-theoretic analysis may be expensive and difficult. But when uncertainties and tradeoffs are dominant features of a decision problem, they cannot be captured in a single heuristic, nor can they be captured explicitly by multiple heuristics (with associated measures of certainty). In combining evidence from rules as if they are modular entities that do not affect the performance of the remaining rules, the implicit assumption is made that these rules are probabilistically independent [27]. Because decision analysis makes explicit the variables on which the success of each heuristic depends, it indicates whether assumptions of modularity are being met. Violating the modularity assumption may have serious implications for system performance [28].

We envision a system where each *situation -> action* heuristic is justified by decision-theoretic knowledge. This will allow the knowledge engineer to estimate the expected gain in system performance when a complete decision analysis is used in place of a simple heuristic. An informed decision can be made between the benefits of the computational economy of heuristics and the possible costs of their computational inaccuracies.

Decision theory represents an important tool that should be considered by expert system builders. Used in conjunction with heuristic techniques, the decision-theoretic approach not only provides a sound basis on which to base knowledge engineering decisions, but also may enhance the ability of a system to explain its reasoning and to solve problems in which the explicit consideration of tradeoffs is essential.

ACKNOWLEDGEMENTS

David Heckerman, Eric Horvitz, and David Wilkins gave helpful comments on an earlier draft of this paper.

REFERENCES

1. Buchanan, B. G., and Shortliffe, E. H., eds.. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, Reading, Mass., 1984.

2. Hayes-Roth, F., Waterman, D. and Lenat, D. (eds.). *Building Expert Systems*. Addison-Wesley, New York, 1983.

3. Yu, V. L., Fagan, L. M., Wraith, S. M., et al. Antimicrobial Selection By A Computer: A Blinded Evaluation By Infectious Disease Experts. *Journal of the American Medical Association 242*, 12 (1979), 1279-1282.

4. Miller, R. A., Pople, H. E., and Myers, J. D. INTERNIST-1, An Experimental Computer-Based Diagnostic Consultant for General Internal Medicine. *New England Journal of Medicine 307*, 8 (1982), 468-476.

5. Hickam, D. H., Shortliffe, E. H., Bischoff, M. B., Scott, A. C., Jacobs, C. D. A Study of the Treatment Advice of a Computer-based Cancer Chemotherapy Protocol Advisor. *Annals of Internal Medicine 103* (1985), 928-936.

6. Clancey, W. J. The Epistemology of a Rule-Based Expert System -- A Framework for Explanation. *Artificial Intelligence 20*, 3 (1983), 215-251.

7. Shortliffe, E. H.. *Computer-Based Medical Consultations: MYCIN*. Elsevier/North Holland, New York, 1976.

8. Lenat, D. B. The Nature of Heuristics. *Artificial Intelligence 19* (1982), 189-219.

9. Smith, R. G., Winston, H. A., Mitchell, T. M., Buchanan, B. G. Representation and Use of Explicit Justifications for Knowledge Base Refinement. Proceedings of the Ninth International Joint Conference on Artificial Intelligence, IJCAI-85, 1985, pp. 673-680.

10. Gaschnig, J. Exactly How Good Are Heuristics?: Toward a Realistic Predictive Theory of Best-First Search. Proceedings of IJCAI-77, International Joint Conference on Artificial Intelligence, 1977, pp. 434-441.

11. Savage, L. J.. *The Foundations of Statistics*. Dover, New York, 1972.

12. Raiffa, H.. *Decision Analysis: Introductory Lectures on Choice Under Uncertainty*. Addison-Wesley, Reading, Mass., 1968.

13. Howard, R. A., Matheson, J. E.. *Readings on the Principles and Applications of Decision Analysis*. Strategic Decisions Group, Menlo Park, CA, 1984. 2nd Edition.

14. Jacobs, W. and Keifer, M. Robot Decisions Based on Maximizing Utility. Proceedings of the Third International Joint Conference on Artificial Intelligence, IJCAI-73, 1973, pp. 402-411.

15. Coles, L. S., Robb, A. M., Sinclair, P. L., Smith, M. H., Sobek, R. R. Decision Analysis for an Experimental Robot with Unreliable Sensors. Proceedings of the Fourth International Joint Conference on Artificial Intelligence, IJCAI-75, 1975, pp. 749-757.

16. Feldman, J. A., Sproull, R. F. Decision Theory and Artificial Intelligence II: The Hungrey Monkey. *Cognitive Science 1* (1975), 158-192.

17. Slagle, J. R. and Hamburger, H. An Expert System for a Resource Allocation Problem. *Communications of the Association for Computing Machinery 28*, 9 (1985), 994-1004.

18. Langlotz, C. P., Fagan, L. M., Shortliffe, E. H. Overcoming Limitations of Artificial Intelligence Planning Techniques. Proceedings of the American Association for Medical Systems and Informatics Congress 1986, Anahiem, California, 1986, pp. 92-96. Also Technical Report KSL-85-25, Knowledge Systems Laboratory, Stanford University

19. Huff, Barbara B. (ed.). *Physicians Desk Reference*. Medical Economics Company, Inc., Oradell, New Jersey, 1985.

20. Langlotz, C., Fagan, L., Tu, S., Williams, J., Sikic, B. ONYX: An Architecture for Planning in Uncertain Environments. Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI-85, Los Angeles, Aug., 1985.

21. Howard, R. A., Matheson, J. E. Influence Diagrams. In *Readings on the Principles and Applications of Decision Analysis*, Howard, R. A., Matheson, J. E., Eds., Strategic Decisions Group, Menlo Park, CA, 1981, ch. 37, , pp. 721-762.

22. Duda, R., Hart, P., and Nilsson, N. Subjective Bayesian Methods for Rule-based Inference Systems. Proceedings 1976 National Computer Conference, AFIPS, 1976, pp. 1075-1082.

23. Tversky, A., Kahneman, D. Judgment Under Uncertainty: Heuristics and Biases. *Science 185* (1974), 1124-1131.

24. Wellman, M. P. Reasoning About Preference Models. Tech. Rept. MIT/LCS/TR-340, Laboratory for Computer Science, Massachusetts Institute of Technology, May, 1985.

25. Pratt, J. W., Raiffa, H., and Schlaifer, R.. *Introduction to Statistical Decision Theory (Preliminary Edition)*. McGraw-Hill, New York, 1965.

26. Doubilet, P., Begg, C. B., Weinstein, M. C., Braun, P., and McNeil, B. J. Probabilistic Sensitivity Analysis Using Monte Carlo Simulation. *Medical Decision Making 5*, 2 (1985), 157-177.

27. Heckerman, D.E. Probabilistic Interpretations for MYCIN's Certainty Factors. In *Uncertainty in Artificial Intelligence*, North Holland, New York, 1986.

28. Horvitz, E. J., and Heckerman, D. E. The Inconsistent Use of Measures of Certainty in Artificial Intelligence Research. In *Uncertainty in Artificial Intelligence*, North Holland, New York, 1986.