



STANFORD RESEARCH INSTITUTE
Menlo Park, California 94025 · U.S.A.

January 1976

SUBJECTIVE BAYESIAN METHODS FOR RULE-BASED
INFERENCE SYSTEMS

by

Richard O. Duda
Peter E. Hart
Nils J. Nilsson

Artificial Intelligence Center

Technical Note 124

SRI Project 4763

The work reported herein was supported by the Advanced Research Projects
Agency of the Department of Defense under Contract DAHC04-75-C-0005.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE JAN 1976		2. REPORT TYPE		3. DATES COVERED 00-01-1976 to 00-01-1976	
4. TITLE AND SUBTITLE Subjective Bayesian Methods for Rule-Based Inference Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SRI International, 333 Ravenswood Avenue, Menlo Park, CA, 94025				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 24	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

SUBJECTIVE BAYESIAN METHODS FOR RULE-BASED
INFERENCE SYSTEMS

R. O. Duda, P. E. Hart, and N. J. Nilsson

ABSTRACT

The general problem of drawing inferences from uncertain or incomplete evidence has invited a variety of technical approaches, some mathematically rigorous and some largely informal and intuitive. Most current inference systems in artificial intelligence have emphasized intuitive methods, because the absence of adequate statistical samples forces a reliance on the subjective judgment of human experts. We describe in this paper a subjective Bayesian inference method that realizes some of the advantages of both formal and informal approaches. Of particular interest are the modifications needed to deal with the inconsistencies usually found in collections of subjective statements.

Index Terms

Inference, Bayes rule, artificial intelligence, production systems, rule-based systems, subjective probability

I INTRODUCTION

One of the characteristics of human reasoning is the ability to form useful judgments from uncertain and incomplete evidence. This ability is not only needed for everyday activities, which people would normally never formalize, but also for tasks such as medical diagnosis or securities analysis, which have been subjected to formal treatment.

Because the general need to form judgments from incomplete data is so widespread, many techniques have been developed to aid or supplant people in this task. Probability theory and statistics provide a powerful framework for dealing with many inference problems [1,2]. In standard approaches, the link between alternative hypotheses and relevant evidence is represented by conditional or joint probabilities that are estimated from statistical samples. If the number of alternative hypotheses and the amount of relevant evidence are not too great, and if the available sample is sufficiently large, then probability and statistics furnish the preferred analytical tools. However, when many kinds of evidence simultaneously bear on an hypothesis, traditional statistical approaches become inappropriate because estimation problems become unmanageable.

Recent work in artificial intelligence has suggested other approaches to the problem of resolving hypotheses on the basis of a mass of uncertain evidence. Among the most attractive are rule-based systems, which use a large body of inference rules, supplied by experts, to provide the knowledge needed to distinguish among competing hypotheses [3-6]. Each inference rule defines the role of a particular set of evidence in

resolving a particular hypothesis. Typically, an ad hoc scoring function is used to combine the effects of collections of uncertain evidence acting through several inference rules on the same hypothesis. Thus, rule-based systems attempt to substitute judgments distilled from long experience for joint probabilities estimated from prohibitively large samples.

Our purpose in this paper is to describe a subjective Bayesian technique that can be used in place of ad hoc scoring functions in rule-based inference systems. Our intent is to retain insofar as possible the well-understood methods of probability theory, introducing only those modifications needed because we are dealing with networks of subjective inference rules. The scope of the paper is limited; we shall not discuss here the more general issues of representation and control that must be faced when designing a complete rule-based inference system.

II FUNDAMENTALS

In a rule-based inference system, the rules are typically of the form

$$\begin{array}{l} \underline{\text{If } E_1 \text{ and } E_2 \text{ and } \dots \text{ and } E_n} \\ \underline{\text{then } H} \end{array}$$

where $E_i (i = 1 \dots n)$ is the i^{th} piece of evidence and H is an hypothesis suggested by the evidence. Each inference rule has a certain strength measured by parameters that will be defined later. For now it suffices to say that the greater the strength, the greater is the power of the evidence to confirm the hypothesis. In most applications, the rules and their strengths are provided by carefully interviewing experts.

The individual pieces of evidence (the E_i) and the hypothesis (H) of a rule are propositional statements. Instead of being either absolutely true or false, the truth values of these propositional statements may be uncertain. In this paper we shall represent these uncertainties by probabilities, so that associated with each propositional statement is a corresponding probability value.

To simplify matters, we shall assume (without loss of generality) that each rule has only a single propositional statement as evidence on its left-hand side. To reduce a conjunction to a single statement, we need a method for computing the joint probability, $P(E_1, \dots, E_n)$ from the individual probabilities $P(E_i)$. Two simple alternatives are to assume independence of the E_i or to use the fuzzy set computation $P(E_1, \dots, E_n) = \min P(E_i)$. More generally, the left-hand side of a rule could contain an arbitrary logical expression, E. The results of this paper do not depend on how the probability of E is computed.


We represent a rule of the form "if E then H" graphically by the following structure:



Here a propositional statement is being represented as a node, and an inference rule is being represented as an arc. A collection of rules about some specific subject area invariably uses the same pieces of evidence to imply several different hypotheses. It also frequently happens that several alternative pieces of evidence imply the same hypothesis. Furthermore, there are often chains of evidences and hypotheses. For these reasons it is natural to represent a collection of rules as a graph structure or inference net.

An example of an inference net is shown in Figure 1. The H_i at the top of the net are alternative hypotheses to be resolved. Each arc entering a node represents an inference rule and has associated with it a strength. Notice that a typical intermediate node like E_5 can play two roles: it provides supporting evidence for the nodes above it (E_2 and E_3), and it acts as an hypothesis to be resolved by evidence below it (E_8 and E_9).

The main problem to be considered in this paper concerns the propagation of probabilities through the net. Suppose for example, that a user of the net provides evidence by deciding that the probability of a node, say E_6 , should be changed from its prior value to some new value. Obviously this should require updating of the probabilities of E_4 and, in turn, E_1 , E_2 , and H_k , and so on. Any mechanism used for propagating probabilities must be able to cope with a number of problems. The rules



SA-4763-1

FIGURE 1 A SIMPLE INFERENCE NET

have uncertainty associated with them, and the evidence provided by a user may be uncertain. These two different kinds of uncertainty must somehow be combined. Multiple evidence typically bears on a single hypothesis, so that some form of independence must usually be assumed. Finally, the rules are provided subjectively by experts, so certain kinds of inconsistencies arise that can seriously jeopardize success. In the following sections we suggest a Bayesian updating scheme that addresses these concerns.

III SUBJECTIVE BAYESIAN UPDATING

Suppose we are given a rule if E, then H. Let us begin with the simplified problem of updating the probability of H given its prior value and given that E is observed to be true. By Bayes rule, we have

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)} \quad (1)$$

For our purposes, a more convenient form of Bayes rule is arrived at by writing the complementary form for the negation of H

$$P(\bar{H}|E) = \frac{P(E|\bar{H})P(\bar{H})}{P(E)} \quad (2)$$

and dividing Eq. (1) by Eq. (2) to obtain

$$\frac{P(H|E)}{P(\bar{H}|E)} = \frac{P(E|H)}{P(E|\bar{H})} \frac{P(H)}{P(\bar{H})} \quad (3)$$

Each of the three terms in this equation has a traditional interpretation. We define the prior odds on H to be

$$O(H) = \frac{P(H)}{P(\bar{H})} = \frac{P(H)}{1 - P(H)} \quad (4)$$

and the posterior odds to be

$$O(H|E) = \frac{P(H|E)}{P(\bar{H}|E)} = \frac{P(H|E)}{1 - P(H|E)} \quad (5)$$

Now the likelihood ratio is defined by

$$\lambda = \frac{P(E|H)}{P(E|\bar{H})} \quad , \quad (6)$$

so Eq. 3 becomes the odds-likelihood formulation of Bayes rule:

$$O(H|E) = \lambda O(H) \quad . \quad (7)$$

This equation tells us how to update the odds on H given the observation of E. For rule-based inference systems, we assume that a human expert has given the rule and has provided the likelihood ratio λ to indicate the "strength" of the rule. A high value of λ ($\lambda \gg 1$) represents, roughly speaking, the fact that E is sufficient for H, since the observation that E is true will transform indifferent prior odds on H into heavy posterior odds in favor of H. Notice, incidentally, that the underlying probabilities can be recovered from their odds by the simple formula

$$P = \frac{O}{O + 1} \quad , \quad (8)$$

so that the odds and the probabilities give exactly the same information.

Suppose now that we wish to update the odds on H given that E is observed to be false. In a strictly analogous fashion, we write

$$O(H|\bar{E}) = \bar{\lambda} O(H) \quad , \quad (9)$$

where we define $\bar{\lambda}$ by

$$\bar{\lambda} = \frac{P(\bar{E}|H)}{P(\bar{E}|\bar{H})} = \frac{1 - P(E|H)}{1 - P(E|\bar{H})} \quad (10)$$

Notice that $\bar{\lambda}$ must also be provided by the human expert; it cannot be derived from λ . A low value of $\bar{\lambda}$, ($0 \leq \bar{\lambda} \ll 1$) represents, roughly speaking, the fact that E is necessary for H, since the observation that E is false will by Eq. 9 transform indifferent prior odds on H into odds heavily against H. Curiously, although λ and $\bar{\lambda}$ must be separately provided by the expert, they are not completely independent of each other. In particular, Eqs. (6) and (9) yield

$$\bar{\lambda} = \frac{1 - \lambda P(E|\bar{H})}{1 - P(E|\bar{H})}, \quad (11)$$

so that, if we exclude the extreme cases of $P(E|\bar{H})$ being either 0 or 1, we see that $\lambda > 1$ implies $\bar{\lambda} < 1$, and $\lambda < 1$ implies $\bar{\lambda} > 1$. Further, we have $\lambda = 1$ if and only if $\bar{\lambda} = 1$. This means that if the expert gives a rule such that the presence of E enhances the odds on H (i.e., $\lambda > 1$), he should also tell us that the absence of E depresses the odds on H (i.e., $\bar{\lambda} < 1$). To some extent, this mathematical requirement does violence to intuition. People who work with rule-based inference systems are commonly told by experts that "The presence of E enhances the odds on H, but the absence of E has no significance." In other words, the expert says that $\lambda > 1$, but $\bar{\lambda} = 1$. Subsequently, we shall suggest some modifications that address this and other problems of inconsistency.

We note in passing that knowledge of both λ and $\bar{\lambda}$ is equivalent to knowledge of both $P(E|H)$ and $P(E|\bar{H})$. Indeed, it follows at once from Eqs. (6) and (10) that

$$P(E|H) = \lambda \frac{1 - \bar{\lambda}}{\lambda - \bar{\lambda}} \quad (12)$$

and

$$P(E|\bar{H}) = \frac{1 - \bar{\lambda}}{\lambda - \bar{\lambda}} \quad (13)$$

Thus, whether the expert should be asked to provide λ and $\bar{\lambda}$, $P(E|H)$ and $P(E|\bar{H})$, or, indeed, some other equivalent information is a psychological rather than a mathematical question [7].

IV UNCERTAIN EVIDENCE AND THE PROBLEM OF PRIOR PROBABILITIES

Having seen how to update the probability of an hypothesis when the evidence is known to be either certainly true or certainly false, let us consider now how updating should proceed when the user of the system is uncertain. We begin by assuming that when a user says "I am 70% certain that E is true," he means that $P(E|\text{relevant observations}) = .7$. We designate by E' the relevant observations that he makes, and simply write $P(E|E')$ for the user's response.

We now need to obtain an expression for $P(H|E')$. Formally,

$$\begin{aligned} P(H|E') &= P(H, E|E') + P(H, \bar{E}|E') \\ &= P(H|E, E')P(E|E') + P(H|\bar{E}, E')P(\bar{E}|E') \end{aligned} \quad (14)$$

We make the reasonable assumption that if we know E to be true (or false), then the observations E' relevant to E provide no further information about H. With this assumption, Eq. (14) becomes

$$P(H|E') = P(H|E)P(E|E') + P(H|\bar{E})P(\bar{E}|E') \quad (15)$$


Here $P(H|E)$ and $P(H|\bar{E})$ are obtained directly from Bayes rule, i.e., from Eq. (7) and Eq. (9), respectively.

If the user is certain that E is true, then $P(H|E') = P(H|E)$. If the user is certain that E is false, then $P(H|E') = P(H|\bar{E})$. In general, Eq. (15) gives $P(H|E')$ as a linear interpolation between these two extreme cases. In particular, note that if $P(E|E') = P(E)$ then $P(H|E') = P(H)$. This has the simple interpretation that if the evidence E' is no better than a priori knowledge, then application of the rule leaves the probability of H unchanged.

In a pure Bayesian formulation, Eq. (15) is the solution to the updating question. In practice, however, there are significant difficulties in using this formulation in an inference net. These difficulties stem from a combination of the classical Bayesian dilemma over prior probabilities and the use of subjective probabilities.

To appreciate the difficulty, consider again a typical pair of nodes E and H embedded in an inference net. It is apparent from Eqs. (7) and (9) that the updating procedure depends on the availability of the prior odds $O(H)$. Thus, although we have not emphasized the point until now, we see that the expert must be depended upon to provide the prior odds as well as λ and $\bar{\lambda}$ when the inference rule is given. On the other hand, recall our earlier observation that E also acts as an hypothesis to be resolved by the nodes below it in the net. Thus, the expert must also provide prior odds on E . If all of these quantities were specified consistently, then the situation would be as represented in Figure 2. The straight line plotted is simply Eq. (15), and shows the interpolation noted above. In particular, note that if the user asserts that $P(E|E') = P(E)$, then the updated probability is $P(H|E') = P(H)$. In other words, if the user provides no new evidence, then the probability of H remains unchanged.

In the practical case, unfortunately, the subjectively obtained prior probabilities are virtually certain to be inconsistent, and the




SA-4763-2

FIGURE 2 IDEALIZED UPDATING OF $P(H|E')$

situation becomes as shown in Figure 3. Note that $P(E)$, the prior probability provided by the expert, is different from $P_c(E)$, the probability consistent with $P(H)$. Here, if the user provides no new evidence--i.e., if $P(E|E') = P(E)$ --then the formal Bayesian updating scheme will substantially change the probability of H from its prior value $P(H)$. Furthermore, for the case shown in Figure 3, if the user asserts that E is true with a probability $P(E|E')$ lying in the interval between $P(E)$ and $P_c(E)$, then the updated probability $P(H|E')$ will be less than $P(H)$. Thus, we have here an example of a rule intended to increase the probability of H if E is found to be true, but which turns out to have the opposite effect. This type of error can be compounded as probabilities are propagated through the net.

Several measures can be taken to correct the unfortunate effects of priors that are inconsistent with inference rules. Since the problem




SA-4763-3

FIGURE 3 INCONSISTENT PRIORS

can be thought of as one of overspecification, one approach would be to relax the specification of whatever quantities are subjectively least certain. For example, if the subjective specification of $P(E)$ were least certain (in the expert's opinion), then we might set $P(E) = P_c(E)$. This approach leads to difficulties because the pair of nodes E and H under consideration are embedded in a large net. For example, in Figure 1, we might be considering node E_2 as the hypothesis H , and node E_5 as the evidence E . If we were to establish a prior probability $P(E_5)$ to be consistent with $P(E_2)$, we would simultaneously make $P(E_5)$ inconsistent with the priors on E_8 and E_9 , which provide supporting evidence for E_5 . Prior probabilities can therefore not be forced into consistency on the basis of the local structure of the inference net; apparently, a more global process--perhaps a relaxation process--would be required.

A second alternative for achieving consistency would be to adjust the linear interpolation function shown in Figure 3. There are several possibilities, one of which is illustrated in Figure 4a. The linear function has been broken into a piecewise linear function at the coordinates of the prior probabilities, forcing consistent updating of the probability of H given E'. Two other possibilities are shown in Figures 4b and 4c. In Figure 4b we have introduced a dead zone over the interval between the specified prior probability P(E) and the consistent prior P_c(E). Intuitively, the argument in support of this consistent interpolation function is that if the user cannot give a response outside this interval, then he is not sufficiently certain of his response to warrant any change in the probability of H. Figure 4c shows another possibility, motivated by the earlier observation that experts often give




*this is really
 $P(H|E, \dots)$
 has about a more
 continuous case*

w/ $P(W, | E_i) \dots$

(a)

SA-4763-4

FIGURE 4 CONSISTENT INTERPOLATION FUNCTIONS




(b)

SA-4763-5

FIGURE 4 CONSISTENT INTERPOLATION FUNCTIONS (Continued)

rules of the form "The presence of E enhances the odds on H, but the absence of E has no significance." By keeping $P(H|E')$ equal to $P(H)$ when $P(E|E')$ is less than $P(E)$ we are effectively allowing the forbidden situation where $\lambda > 1$ and $\bar{\lambda} = 1$. In effect, this is equivalent to the method illustrated in Figure 4a under the assumption that $P(H|\bar{E}) = P(H)$.

It is interesting to compare these modifications with the procedure used by Shortliffe to handle uncertain evidence in the MYCIN system [4,5]. While the nonlinear equations that result from use of Shortliffe's version of confirmation theory prevent a general comparison, it is possible to express his procedure in our terms for the special case of a single rule. The result for the case in which the presence of E supports H is shown in Figure 5. ~~Clearly~~ the solution is identical to that of




(c)

SA-4763-6

FIGURE 4 CONSISTENT INTERPOLATION FUNCTIONS (Concluded)

Figure 4c except for the interval from $P(E)$ to $P_t(E)$ within which Shortliffe's solution maintains $P(H|E')$ at the a priori value $P(H)$.

The graphical representations in Figures 2 through 4 provide a nice vehicle for visualizing the discrepancies between formal and subjective Bayesian updating, and make it easy to invent other alternatives for reconciling inconsistencies. For completeness, the Appendix contains the easily computable algebraic representations of these functions, and also treats the complementary case in which the straight line given by Eq. (15) has a negative slope (the case in which $\lambda < \bar{\lambda}$). In a small experimental system, the function shown in Figure 4a has given satisfactory preliminary results [8].



SA-4763-7

FIGURE 5 THE INTERPOLATION FUNCTION USED IN THE MYCIN SYSTEM
 $P_t(E) = P(E) + t[1 - P(E)]$. Typically, $t = 0.2$.

V THE USE OF MULTIPLE EVIDENCE

We turn now to the more general updating problem in which several rules of the form $E_1 \rightarrow H, \dots, E_n \rightarrow H$ all concern the same hypothesis H .^{*} Since most nodes in actual inference nets have several incoming arcs, this is the case of greatest practical interest. In order to gain some insight about how multiple evidence should be used to update H when the evidence is uncertain and the priors are inconsistent, let us first consider briefly how updating would formally proceed in simpler cases.

Suppose the i^{th} inference rule has associated with it the usual two quantities λ_i and $\bar{\lambda}_i$. For a first simple case, how should H be updated

* This should not be confused with the conjunctive premise mentioned earlier.

when all the E_i have been observed to be certainly true? This case is analogous to the case summarized by Eq. (7). Under the assumption that the pieces of evidence are conditionally independent (i.e., that $P(E_1, \dots, E_n | H) = \prod_{i=1}^n P(E_i | H)$ and that $P(E_1, \dots, E_n | \bar{H}) = \prod_{i=1}^n P(E_i | \bar{H})$), it is not difficult to reach an analogous answer. Specifically, the odds on H are updated by the expression

$$O(H | E_1, \dots, E_n) = \left[\prod_{i=1}^n \lambda_i \right] O(H) \quad , \quad (16)$$

where

$$\lambda_i = \frac{P(E_i | H)}{P(E_i | \bar{H})} \quad . \quad (17)$$

Similarly, if all the evidence is observed to be certainly false, we can under conditional independence assumptions again factor the joint likelihood ratio to obtain

$$O(H | \bar{E}_1, \dots, \bar{E}_n) = \left[\prod_{i=1}^n \bar{\lambda}_i \right] O(H) \quad . \quad (18)$$

Now let us consider the general case of uncertain evidence and inconsistent prior probabilities. We already know that the posterior odds $O(H | E'_i)$ given a single observation E'_i can be computed using updating functions like the ones shown in Figure 4. We can therefore define, for a single inference rule, an effective likelihood ratio λ'_i by

$$\lambda'_i \triangleq \frac{O(H | E'_i)}{O(H)} \quad . \quad (19)$$

By making the assumption now that the E'_i are independent, we can obtain for the general case an expression similar to the simple updating formulas given by Eqs. (16) and (18):

$$O(H|E'_1, \dots, E'_n) = \left[\prod_{i=1}^n \lambda'_i \right] O(H) \quad (16)$$

To use this expression in an inference net system, we simply store with each node its prior odds (or probability), and store with each incoming arc an effective likelihood ratio λ'_i . Whenever a piece of evidence provided by the user causes $P(E_i|E'_i)$ to be updated, a new effective likelihood ratio is computed and the posterior odds in favor of H is computed using Eq. (20). This procedure has the following consequences:

- (1) If no evidence is obtained for a rule, then it will retain an initial effective likelihood ratio of unity, since prior and "posterior" odds are the same.
- (2) The order in which evidence is obtained and rules are applied does not affect the final posterior probabilities.
- (3) The same rule can be used repeatedly, with the same or different values for the probability of the evidence. In particular, if a user changes his mind and modifies an earlier assertion, the new assertion will correctly "undo" any effects of earlier statements.

> is this true w/ a threshold?

VI CONCLUSIONS

The probability updating procedure presented here has several points to recommend it. It accepts subjective information that can readily be obtained from experts. The two conditional probabilities, $P(E|H)$ and $P(E|\bar{H})$, that determine the strength of an inference rule typically are intuitively meaningful measures, and the procedure is tolerant of the inevitable inconsistencies in subjective expert information. The basis in probability theory of our procedure provides a useful theoretical foundation for calculating the effects of uncertain evidence. One value of theory is that it makes us explicitly aware of certain underlying assumptions about such matters as conditional independence, prior

probabilities, and inconsistent information. Finally, our procedure is straightforward computationally and can be readily implemented in inference net systems.

There are, however, some questions that remain to be dealt with. If the network contains multiple paths linking a given piece of evidence to the same hypothesis, the independence assumption is obviously violated. It is important to settle on a reasonable (if ad hoc) modification of our basic procedure that behaves appropriately in such situations. (A more extreme complication would involve being able to avoid the circular reasoning implied by inference nets with loops.)

There are sometimes cases where some of the nodes in an inference net are related by a constraint not expressed in any given rule. For example, a subset of hypotheses may be mutually exclusive and exhaustive, in which case their probabilities must always sum to one, regardless of their individual values. Such a constraint may be inconsistent with the associated rule strengths given us by the experts. Perhaps a simple expedient, such as renormalization of probability values, can be justified in this case.

We have not addressed here at all issues of inference net control strategy: for example, which hypotheses should be pursued and which evidence should be sought at any step. The answers to these sorts of questions may be heavily dependent on the particular application. Another global question concerns rules containing logical statements that may include quantifiers and variables. But in whatever way these questions are answered, the basic updating procedure presented here would appear to be a useful component of rule-based inference systems.

ACKNOWLEDGMENTS

We have benefitted from the comments of many of our colleagues, but would like particularly to acknowledge the contributions of Georgia Sutherland at SRI, and the stimulating and helpful discussions with E. H. Shortliffe, Bruce Buchanan, Randall Davis, and Dana Ludwig at Stanford University. This work was supported by the Advanced Projects Research Agency under Contract DAHC04-75-C-0005.

APPENDIX

Complete analytical expressions giving $P(H|E')$ as a piecewise linear function of $P(E|E')$ are given in this Appendix. These expressions correspond to the three graphical representations illustrated in Figure 4. The simplest expression corresponds to Figure 4a:

$$P(H|E') = \begin{cases} P(H|\bar{E}) + \frac{P(E|E')}{P(E)} [P(H) - P(H|\bar{E})] & 0 \leq P(E|E') \leq P(E) \\ \frac{P(H) - P(H|E)P(E)}{1 - P(E)} + P(E|E') \frac{P(H|E) - P(H)}{1 - P(E)} & P(E) \leq P(E|E') \leq 1 \end{cases} \quad (A1)$$

Here it is important to note that the four quantities $P(H)$, $P(E)$, $P(H|E)$, and $P(H|\bar{E})$ are assumed to be estimates obtained from experts. Were the true probabilities to be used in this formula, it would reduce at once to the linear expression given by Eq. (15). The estimates of $P(H|E)$ and $P(H|\bar{E})$ might be obtained directly from an expert, but would more often be obtained through Bayes rule [Eqs. (7) and (9), respectively]. To be explicit,

$$P(H|E) = \frac{P(E|H)P(H)}{[P(E|H) - P(E|\bar{H})]P(H) + P(E|\bar{H})} = \frac{\lambda P(H)}{(\lambda - 1)P(H) + 1} \quad (A2)$$

and

$$P(H|\bar{E}) = \frac{[1 - P(E|H)]P(H)}{[P(E|\bar{H}) - P(E|H)]P(H) + 1 - P(E|\bar{H})} = \frac{\bar{\lambda}P(H)}{(\bar{\lambda} - 1)P(H) + 1} \quad (A3)$$

To obtain the equations for Figure 4b, we define $P_c(E)$ by

$$P_c(E) = \frac{P(H) - P(H|\bar{E})}{P(H|E) - P(H|\bar{E})} \quad (A4)$$

In general, this quantity will differ from the $P(E)$ value supplied by the expert. For Figure 4b we must distinguish between the two cases $P(E) \leq P_c(E)$ and $P(E) > P_c(E)$. The equations are as follows:

Case 1: $P(E) \leq P_c(E)$

$$P(H|E') = \begin{cases} P(H|\bar{E}) + \frac{P(E|E')}{P(E)} [P(H) - P(H|\bar{E})] & 0 \leq P(E|E') \leq P(E) \\ P(H) & P(E) \leq P(E|E') \leq P_c(E) \\ P(H|\bar{E}) + P(E|E') [P(H|E) - P(H|\bar{E})] & P_c(E) \leq P(E|E') \leq 1 \end{cases} \quad (A5)$$

Case 2: $P(E) > P_c(E)$

$$P(H|E') = \begin{cases} P(H|\bar{E}) + P(E|E') [P(H|E) - P(H|\bar{E})] & 0 \leq P(E|E') \leq P_c(E) \\ P(H) & P_c(E) \leq P(E|E') \leq P(E) \\ \frac{P(H) - P(H|E)P(E)}{1 - P(E)} + P(E|E') \frac{P(H|E) - P(H)}{1 - P(E)} & P(E) \leq P(E|E') \leq 1 \end{cases} \quad (A6)$$

Finally, there are also two cases to be distinguished for Figure 4c. The first case corresponds to assuming that $P(H|\bar{E}) \approx P(H)$, so that $P_c(E) \approx 0$. The second case corresponds to assuming that $P(H|E) \approx P(H)$, so that $P_c(E) \approx 1$. In effect, these cases correspond to the rules $E \overset{\lambda}{\Delta} H$ and $\bar{E} \overset{\bar{\lambda}}{\Delta} H$ taken separately. The corresponding equations are special cases of Eqs. (A5) and (A6):

Case 1: $E \xrightarrow{\lambda} H$

$$P(H|E') = \begin{cases} P(H) & 0 \leq P(E|E') \leq P(E) \\ \frac{P(H) - P(H|E)P(E)}{1 - P(E)} + P(E|E') \frac{P(H|E) - P(H)}{1 - P(E)} & P(E) \leq P(E|E') \leq 1 \end{cases} \quad (A7)$$

Case 2: $\bar{E} \xrightarrow{\bar{\lambda}} H$

$$P(H|E') = \begin{cases} P(H|\bar{E}) + \frac{P(E|E')}{P(E)} [P(H) - P(H|\bar{E})] & 0 \leq P(E|E') \leq P(E) \\ P(H) & P(E) \leq P(E|E') \leq 1 \end{cases} \quad (A8)$$

Ordinarily one would view this as a simplified approximation that is useful when one of the two likelihood ratios is dominant. However, it is interesting to observe that if both λ and $\bar{\lambda}$ are significant and if the two separate rules $E \xrightarrow{\lambda} H$ and $\bar{E} \xrightarrow{\bar{\lambda}} H$ are treated as if E and \bar{E} were statistically independent, then Eqs. (A7) and (A8) yield the same result as Eq. (A1). This follows from the fact that when $P(H|E') = P(H)$ we have $O(H|E') = O(H)$, so that Eq. (19) yields $\lambda' = 1$. Thus, if $0 \leq P(E|E') \leq P(E)$ only the rule $\bar{E} \xrightarrow{\bar{\lambda}} H$ contributes to $P(H|E')$, while if $P(E) \leq P(E|E') \leq 1$ only the rule $E \xrightarrow{\lambda} H$ contributes to $P(H|E')$, the contributions being exactly those given in Eq. (A1).

REFERENCES

1. G. Hadley, Introduction to Probability and Statistical Decision Theory (Holden-Day, San Francisco, California, 1967).
2. H. Raiffa, Decision Analysis (Addison-Wesley, New York, New York, 1968).
3. D. A. Waterman, "Generalization Learning Techniques for Automating the Learning of Heuristics," Artificial Intelligence, Vol. 1, pp. 121-170 (Spring 1970).

4. E. H. Shortliffe, "MYCIN: A Rule-Based Computer Program for Advising Physicians Regarding Antimicrobial Therapy Selection," Stanford Artificial Intelligence Laboratory Memo AIM-251, Stanford University, Stanford, California (October 1974).
5. E. H. Shortliffe and B. G. Buchanan, "A Model of Inexact Reasoning in Medicine," Mathematical Biosciences, Vol. 23, pp. 351-379 (1975).
6. R. Davis and J. King, "An Overview of Production Systems," in Machine Representations of Knowledge (D. Reidel Publishing Co; forthcoming).
7. D. H. Gustafson et al., "Wisconsin Computer Aided Medical Diagnosis Project--Progress Report," in Computer Diagnosis and Diagnostic Methods, pp. 255-278, J. A. Jacquez, ed. (Charles C. Thomas, Springfield, Illinois, 1972).
8. G. Sutherland, "Implementation of Inference Nets--II," Technical Note 122, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California (January 1976).