

EXTENDED PLAUSIBLE INFERENCE

Leonard Friedman

Jet Propulsion Laboratory, C.I.T.
Pasadena, CA 91109

ABSTRACT

A generalized theory of plausible inference has been developed, extended to arbitrary expressions of propositional calculus from Shortliffe and Buchanan's original MYCIN formulation. The theory represents uncertainty of belief, and invokes four rules of inference, instead of the two of standard logic. Areas of application include diagnostic problems and deciding between alternative hypotheses. The theory has been implemented in the PI system. The intended area of application for PI is the trouble-shooting of a failed spacecraft, and the solution to a simplified problem drawn from this area is described.

FOREWORD

After this paper was submitted for review, the work of two mathematicians, Arthur Dempster and Glenn Shafer, was brought to my attention. Shafer has published a book, 'A Mathematical Theory of Evidence', which establishes a rigorously founded theory of reasoning under uncertainty [1]. Plausible inference is based on the uncertainty formulas devised by Shortliffe and Buchanan, who took pains to point out that their formulas were ad hoc [2] or [3]. Comparison with Shafer's formulas reveals that the formulas employed in plausible inference must be revised to be correct. The principal errors are of two kinds: neglect of conflicting evidence, for which the conjunction is false with certainty; and neglect of set inclusion (whereby some assertions either may or may not be arranged as subsets of others).

In particular, Theorem 10.4, on page 224 of the paperback edition of Shafer's book states that the minimum rule for conjunctions (support for a conjunction is the minimum of the supports for its arguments) applies if and only if the assertions that are the arguments are 'consonant'; i.e., that they can be arranged by some criterion for set inclusion and accordingly do not support conflict in the evidence. This, of course, also applies to the maximum rule for disjunctions. Shafer gives the correct rules for heterogeneous and conflicting

* This paper represents the results of one phase of research carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

evidence in paragraphs 4 and 5 of chapter 4. Heterogeneous evidence refers to conjunctions which may be true, whose arguments cannot be arranged in subsets, and which may support conflict in the evidence. Support for this type of conjunct is given by the product of the support for the arguments. Conflicting evidence refers to conjuncts whose support is guaranteed to be false with certainty (support = 0). Thus we discover that the principal impact of the corrections to plausible inference is on the formulas for calculating the support for conjunctions and disjunctions, and the knowledge required as to when to apply which formula.

Happily, the conceptual framework of plausible inference survives intact, as do many of the formulas. For example, the formula for 'credibility summation' is correct for consonant and heterogeneous evidence. Dempster's rule for orthogonal summation reduces to the same result for those cases. A different summation formula results when evidence is guaranteed to be conflicting. We retain the assumption of directed propagation of support in implications, whereby conditional support is proportional to a 'relevance factor'. These calculations, based on the relevance factors or ΔI 's, are not considered by Shafer. Various inconsistencies are removed by employing Shafer's formulas. In particular, the law of the excluded middle always holds, as well as normalization, so that belief in a defined 'frame of discernment' or universe of discourse always adds up to 1, and the conjunction of support for and against an assertion is always certainly false.

Note also that two other papers in these proceedings rely on Dempster and Shafer's work. See [4] and [5].


BACKGROUND

One of the basic assumptions of standard logic is that the truth value of predicates is known with certainty. This assumption is justified for some domains of reasoning, but is not valid when, for example, we are confronted by problems of diagnosis or hypothesis generation in scientific experiments. In such domains there may be conflicting evidence for, or against particular assertions and a human, when reasoning about them uses measures of belief other than true or false to characterize his uncertainty.

Logicians and philosophers have studied this phenomenon for many years and have proposed various criteria as formal models, particularly for confirming evidence. The mathematician, George Polya, wrote a book about the modes of reasoning employed by fellow mathematicians, and emphasized the importance of uncertainty even in mathematical reasoning. He suggested that humans used non-standard rules of inference that included confirmation and denial in addition to ponens and tollens [6]. Recently Zadeh and his school have developed the concept of fuzzy logic which deals extensively with uncertainty [7]. An approach which has been useful in a limited domain of medical diagnosis was formulated by Shortliffe and Buchanan for their AI system, MYCIN. MYCIN is a production system and reasons only in a confirmation mode for a particular subset of logical expressions in the propositional calculus.

This paper presents a generalization of Shortliffe and Buchanan's theory to the four modes of reasoning suggested by Polya, and to arbitrary expressions of the propositional calculus. The resultant theory of reasoning with uncertain beliefs is called plausible inference and has been implemented as the PI system. PI has been applied to the trouble-shooting of a spacecraft after the execution of an experiment has failed, and this example's operation will be described briefly here.

Shortliffe and Buchanan (hereafter referred to as S and B) give an extensive discussion of their choice of measures of uncertainty as functions of conditional probabilities. In particular, they point out that humans reasoning about uncertain assertions do not employ evidence FOR an assertion to measure their disbelief. Independent evidence is required in normal human reasoning to bolster disbelief in the assertion. As a result their model and ours treat measures of belief and disbelief as independent quantities.




MYCIN/PROSPECTOR TREE
Figure 1

This is unlike conditional probability, which asserts $P(\sim Z|A) = 1 - P(Z|A)$. It is evident then that these measures are NOT probabilities, however tempting it may be to regard them as such. We shall not duplicate S and B's justification of the theoretical foundations of the measures we are using which is fundamentally based on its success in modelling human reasoning. Schefe has also investigated the foundations of this model, and concludes that it is valid [8]. In Schefe's opinion, Zadeh's 'Linguistic Modelling' is 'inadequate' (his wording).

We shall begin by comparing the basic reasoning methods employed by or available to PI with those used by the MYCIN and PROSPECTOR systems. Then we shall describe the example actually running. Finally, the details of the underlying plausible inference theory will be explained.

COMPARISONS AND POTENTIALS

MYCIN and PROSPECTOR employ similar inference strategies to reach conclusions. PROSPECTOR is described in [9]. In both systems a tree structure of evidential assertions supporting various hypotheses converges to several root conclusions. (See Fig. 1) The reasoning is essentially one directional and in the confirmation mode, with the evidence given possibly supporting more than one conclusion. The choice of conclusion is then made on the basis of degree of certainty or subjective probability assigned to each conclusion that is supported by the evidence, with the most certain being chosen. Degree of certainty is propagated to the conclusions by setting up a weighting system of numbers assigned to each arc linking the nodes. The arc numbers are proportionality factors that determine the amount of confidence transmitted from evidence towards conclusions. Once confidence in the evidence has been assigned, various combining rules determine the propagation of certainty toward conclusions,



PI TREE FOR IMPLEMENTED EXAMPLE
Figure 2

with the system working in the confirmation mode.

The PI working example has a tree structure shown in Fig. 2. The nodes converge to a goal state which is the root of the tree. In this example, we assume that the nodes constitute a plan of physical processes and conditions leading to the goal state, and that, after execution of the plan, the goal state is not achieved; i.e., it is false. Thus we are seeking to isolate the cause of the failure in one of the leaves of the tree. We proceed from the root of the tree toward the leaves, attempting to isolate the failed leaf by pruning off those parts of the tree closest to the root that are still working; i.e., that are true. Once again we are proceeding in essentially one direction, in a rather rigid manner as compared with a human, but in this example PI employs two modes, tollens and confirmation.

PI also requires the same type of numbers assigned to its arcs, and I have given them the name 'relevance factors'. In both the MYCIN and PROSPECTOR systems they are determined by the subjective estimates of human experts in the knowledge domain, and are not altered except for system debugging purposes. By contrast, for the type of solution method described above, the PI system has been provided with a simple heuristic that permits it to estimate how the value of the relevance factors should be altered in a given neighborhood, on the basis of evidence it receives. This ability enables it to work on problems for which the relevance factors are not known in advance, and prepares the way for embedding in a learning system.

Clearly we are not exploiting the full power of the PI system by using only the solution methods illustrated by Fig. 2. A far more interesting and efficient approach is described by Fig 3. We presume that for any realistic problem the tree of Fig. 2, describing the states and actions, will be very large. Fig. 3 shows only a small neighborhood of that tree. What a human very

often does, when confronted with such a problem, is to select, on the basis of (usually unconscious) criteria and inferences, a specific neighborhood of the web of relations, and test the function represented by that neighborhood.

An augmented PI might be able to do the same thing by using a set of criteria stored in a separate knowledge base. Examples of such rules are, 'What was the last thing altered since this plan worked correctly?' or 'What is the cheapest test that can be performed verifying the nodes of the tree?'. Such a body of knowledge would form a meta-rule domain of the kind described by Davis, and the same PI inference engine that reasons about spacecraft experiments could be used to reason about where to choose the next neighborhood for testing[10]. In fact, in this case all the nodes of the tree are candidates for control selection, making his approach imperative.

Let us assume that either an augmented PI system has selected a node for checking as shown in Fig. 3, or simply that a human has suggested it. If a test is then made that determines the truth or falsity of that node, a degree of certainty can be assigned to that node. The PI system already has the capacity to propagate that change in certainty in both directions, thus providing much more information on the choice of the next test to perform.


AN EXAMPLE OF FAULT DIAGNOSIS

The specific example implemented is intended to exercise some of the belief propagation processes of plausible inference. It also demonstrates the feasibility of writing an algorithm for estimating relevance factors, which measure the strength of binding between antecedents and consequents as evidence is adduced.

This is an unrealistically small problem. (See [11] for a detailed description.) The pruning methods employed are clearly too costly for large plans so long as the method of solution is that of making every test from the root towards the leaves as in Fig 2. The example itself is taken from the application domain being worked on at JPL, and uses as its knowledge base a model of normal operation of a spacecraft generated by another operating module, the JPL planner, for commanding the spacecraft.


The personnel monitoring the spacecraft are rarely experts in all of the complex sub-systems that can fail. Therefore it would be most useful if these already available plans of normal operation could be made to serve as guides for focussing attention on specific subsystems, a kind of spacecraft first-aid. For in-depth fault diagnosis, detailed models of failure modes would be required, just as doctors and AI researchers have built up models of disease to guide automatic diagnosis systems.

The planner is the part of the JPL AI system which generates a plan for achieving a specified goal state. The actions leading to



PI DEBUGGING IN A VERY LARGE TREE
Figure 3


sub-goal states are limited to the repertoire of commands of the spacecraft computer, so that when these actions are stripped out after a plan is generated, they constitute a valid program for the computer. Thus the planner is a simple automatic code generator. It uses a representation of physical causality borrowed from Chuck Rieger's Common Sense Algorithm work [13]. Figure 4 shows the representation for an action causing a state provided that the pre-conditions, state1 and state2, are first true.



"ACTION" CAUSES "STATE"
PRE-CONDITIONS ARE STATE 1 AND STATE 2

REPRESENTATION OF CAUSALITY
Figure 4

In Figure 5 a simplified plan is shown for transmitting a signal from the spacecraft to earth. The action of 'Ground Antenna Receiving' is required, plus the pre-conditions that (1) the spacecraft is pointed correctly at the earth, and (2) a signal was transmitted by the spacecraft. The pre-conditions, in turn, require actions and pre-conditions to make them true, etc. to the desired level of representation of the physical processes involved.



PLAN FOR TRANSMITTING FROM SPACECRAFT TO EARTH
Figure 5

To reason in a logically correct manner about causal relations of this type requires a special function, the CAUSAL specialist, that can take into account the temporal relations, verifying that the pre-conditions were made true before the action process was started. The JPL planner has such a function, but a similar one has not yet been incorporated in PI. Instead, the logical relations have been approximated by representing the action

and pre-conditions as arguments of a conjunction implying the subgoal or goal state. If the goal state is the assertion A, it may be represented by $[(Z1 \wedge \dots \wedge Zn) \rightarrow A]$, where $Z1, \dots, Zn$ are the arguments representing action and pre-conditions.

In operation, we assume that a user would have first generated the plan to obtain a command sequence that would run the spacecraft computer to achieve the desired goal. If, after commanding the spacecraft to execute the plan, a failure were detected, that information would be entered into the data base as ['Signal Received' is false.]. This results in a propagation of negative credibility to all the conditions (actions and pre-conditions) on which the failed state depended. We also propagate a 'marker' truth value of PF, standing for possibly false, to these conditions. As far as inference is concerned, PF behaves exactly like false, except that it cannot override an F or T state. This is because it is based on a priori or presumptive evidence only, whereas F or T are based on direct evidence.

The PI system models event succession by using PLANNER-like context layers. Thus it can deal with non-monotonic logic, where a given state can be true or false with varying degrees of certainty as successive actions are taken. A non-monotonic inference rule has been implemented in the following manner. Suppose a goal state depends on n conditions being true to be attained. If, after executing the process, the goal state is not observed, then we conclude that one of the n conditions must be false, and all are possibly false. If, after performing appropriate tests, a human gathers evidence that one of the conditions is true and enters that fact, it may be eliminated from consideration as a cause of failure, and belief in the falsity of the remaining conditions should increase.

If all but one condition is eliminated in this way, PI concludes that the remaining condition must be false. If this condition depends in turn on others, the same process may be applied again to narrow down the suspected fault to a leaf of the tree. Of course, if at some stage testing establishes that all the conditions on which a given failed condition directly depend are true, then PI concludes that the failure condition is not represented in the plan.

PI proceeds by reasoning in the manner just described in either of two computing modes. The first mode duplicates the MYCIN and PROSPECTOR approach in that a priori relevance factors are entered in advance. They constitute a human estimate about the certainty of the falseness of the arguments of the antecedent conjunction, given that the consequence is false. The second is a default mode, in which there is no such information. In default, the program counts the number of arguments and assigns a relevance factor of $1/n$ for n arguments. For either mode, A is either false or possibly false and the Zi are possibly false. If a subset $\{Zj\}$ of arguments of the conjunction are then shown to be true, the remaining arguments form a new subset $\{Zi\}$ that are

possibly false with increased certainty. The relevance factors, Δ_i 's in our notation, are then given by

$$\Delta_{\text{new}} = \left[\Delta_{\text{old}} / \sum_i \Delta_i \right] * \min \{C(Z_j)\}.$$

where $C(Z_j)$ is the degree of certainty or credibility of Z_j . The simple ratio is reduced to the degree that we lack confidence in the beliefs of the factors testing true. If there are three factors, then in the default mode the Δ_i 's are all $1/3$ without test knowledge. If one of the factors has been shown to be true with certainty, we get $\Delta_{\text{new}} = 1/3 / 2/3$, or $\Delta_{\text{new}} = 1/2$, as it should be. Figure 5 shows a priori relevance factors entered. For this mode, if the Ground Antenna Receiving tested true, we would get $\Delta_{\text{new}} = .3$ or $.7$, for the remaining PF factors, reflecting the human judgements accurately.

EXTENDED PLAUSIBLE INFERENCE

One of the basic concepts of plausible inference is credibility. This is a measure of the certainty that a given assertion A is true or false, and is denoted $C(A)$. It is a number lying between -1 and $+1$. -1 denotes false with certainty, $+1$ represents true with certainty, and 0 stands for no information. Intermediate values stand for various degrees of certainty. (This quantity is called the Certainty Factor by S and B. I have introduced a new notation that clarifies the role of the parameters, and permits representing many equations by a single equation. For those familiar with Shortliffe and Buchanan's notation, Appendix I defines the equivalent terms.)

In standard propositional calculus, a set of assertions may be linked by connectives such as AND, OR or IMPLIES. If any assertion linked to others by such connectives changes truth value, this 'disturbance' may be propagated to the linked assertions by well-known rules. The calculus of plausible inference consists of rules to determine the propagation of credibility to the linked assertions, and permits propagation of disturbances in directions forbidden in propositional calculus.

There are two distinctively different ways in which this propagation may take place. The first applies only to the antecedents and consequents of implications and is called directed propagation. The second way is deduced propagation; the assertions to be changed must be deduced from the circumstances. Antecedents or consequents transmit all or a fraction of their credibility to their partners, and only their partners, when they change credibility. If we regard assertions linked by implication as nodes in an implication space, then we have propagation along the arcs connecting linked nodes. (See Fig. 6) The amount propagated depends on their own credibility reduced by a factor called the relevance factor, Δ_i , defined below. The relevance factor may be regarded as a directed quantity, somewhat like a vector, while credibility is like a scalar. Fig. 6 shows that there are at least four Δ_i 's for every node pair linked by

implication.

Deduced propagation applies to the credibilities associated with the connectives AND, OR, and IMPLIES themselves rather than the antecedents or consequents of implications. The collective credibilities of such expressions are given by the maximum or minimum credibility of one of their arguments. This method of assigning credibility was suggested by Zadeh [7]. Deduced propagation also applies to a case such as an OR expression that is defined true. If all but one of its arguments is known to be false, the remaining argument must have a positive credibility regardless of which one it might be. This difference in the way credibility is propagated models our intuitive notion that the believability of the members of an implication is linked; whereas, the arguments of a conjunction or disjunction are treated as independently believable factors.

As pointed out above, we desire to model the independence of confirming and denying evidence, so we need separate measures of amount of belief for or against an assertion 'A'. These are denoted $I(A)$ and $D(A)$, standing for the sum of all increments to belief in A and the sum of all decrements to belief in A. The summation is non-linear and will be defined below. Also,

$$I(A) = D(\sim A).$$

In words, the belief in A true equals the disbelief in A false. We define


$$C(A) = I(A) - D(A).$$

Credibility can only be determined after having calculated I and D. Also, by interchanging increments and decrements, we get

$$C(\sim A) = -C(A).$$

We now introduce a number Δ_i to represent the increment or decrement propagated between

- ARCS ARE LINKS BETWEEN ANTECEDENTS AND CONSEQUENCES.
- NODES ARE ASSERTIONS.



SEMANTIC NET
Figure 6

antecedent and consequent of an implication. Our notation is:

$\Delta(Z|A)$, the increment or decrement propagated from A to Z,
where $0 \leq \Delta(Z|A) \leq 1$.

There are three cases to distinguish in computing the transfer of belief between antecedents and consequents in an implication. Case 1 applies when neither the antecedent nor consequent is an assertion with multiple arguments; i.e., neither A nor Z is an assertion involving AND, OR, or IMPLIES. Case 2 applies when one assertion has multiple arguments and one does not. Case 3 applies when both antecedent and consequent have multiple arguments.

Case 1. Neither antecedent nor consequent has multiple arguments.

The fundamental rule of plausible inference is:

$$\Delta(Z|A) = \Delta_1(Z|A) * |C(A)|,$$

where Δ_1 is the value of Δ when $|C(A)|$ is 1,
A is the assertion with changed credibility.
Z is the target assertion.
Propagation is from A to Z.
TV(A) is truth-value of A.

If $A \rightarrow Z$, and $TV(A) = T$;
 $\Delta(Z|A)$ is added to $I(Z)$; Ponens

$TV(A) = F$;
 $\Delta(Z|A)$ is added to $D(Z)$; Denial.

If $Z \rightarrow A$, and $TV(A) = T$;
 $\Delta(Z|A)$ is added to $I(Z)$; Confirmation

$TV(A) = F$;
 $\Delta(Z|A)$ is added to $D(Z)$; Tollens.

Note that unlike propositional calculus, in plausible inference, belief is propagated in the Denial and Confirmation modes. Note also that Ponens and Confirmation are equivalent, as are Denial and Tollens. This equivalence suggests that plausible reasoning could be symmetric; i.e., for a given knowledge base, if every antecedent and consequent were interchanged, together with their Δ_1 's, the same inferences would be made given the same evidence. Indifference to the direction of reasoning would model what appears to be a similar human capability to reason with antecedents and consequents interchanged. Initially, as this system was developed from propositional calculus, exclusive use of a non-symmetric connective, implication, for directed propagation of credibility concealed this possibility [13].

Another connective is available for representing directed propagation that is symmetric, however. Equivalence differs from implication only with respect to one truth value. When antecedent is false and consequent is true, implication is true and equivalence is false. There are situations where equivalence might be a more suitable connective to employ in the data

base, in which case symmetric reasoning would apply. Such situations bear investigation.

Note also that in plausible reasoning the connectives have a different meaning than that conveyed in propositional calculus. Thus, in plausible reasoning, use of the equivalence connective would only imply a connection in belief, not logical equivalence. Only in the limiting case where plausible inference degenerates to propositional calculus does logical equivalence hold. This would occur when the relevance factors were all equal to 1 and the credibilities were limited to 1 or -1. In what follows we shall use the implication connective because of its familiarity, but corresponding formulas for equivalence are easily derived.

Case 2. One argument of the implication has multiple arguments. The other argument does not.

The logical statements of greatest interest are those involving the AND and OR connectives. To avoid the confusion of too much detail, we shall give the rules of Case 2 only for those connectives. The corresponding rules for antecedents and consequents which are themselves implications is given in Appendix II. All implication rules are derived from the AND/OR rules in any event.

Before giving the details of Case 2, we note that the following assumptions are made for plausible inference. If $I(Z) = 1$, then evidence for disbelief is disregarded. If $D(Z) = 1$, then evidence for belief is disregarded. If both equal 1, then we have contradiction. The calculation of credibility is the same regardless of the order of application of evidence (commutativity).

Case 2a. $[(\text{CONN } A_1 \dots A_n) \rightarrow Z]$,
or $[Z \rightarrow (\text{CONN } A_1 \dots A_n)]$.

where $(\text{CONN } A_1 \dots A_n)$ is the source with changed credibility,

Z is the target assertion,
CONN is one of AND, OR, NOT AND, NOT OR.

Case 2a-1.

If CONN is AND or NOT OR with $TV(\text{CONN}) = T$,
OR or NOT AND with $TV(\text{CONN}) = F$,

then after calculation of the collective credibility of the source expression, case 2a-1 reduces to Case 1 with $A = (\text{CONN } A_1 \dots A_n)$. The reduction is done according to the following rules:

Let the set $A_1 \dots A_n$ be represented by $\{A_i\}$.

For CONN = AND or NOT OR:

$$I(A) = \min \{ I(A_i) \},$$

$$D(A) = \max \{ D(A_i) \},$$

where $\sim A_i$ is substituted for A_i in NOT OR.

For CONN = OR or NOT AND:

$$I(A) = \max \{ I(A_i) \},$$

$$D(A) = \min \{ D(A_i) \},$$

where $\sim A_i$ is substituted for A_i in NOT AND.

Therefore, $\Delta(Z|A) = \Delta_1(Z|A) * |C(A)|$, added to $I(Z)$ for AND, NOT OR true, and to $D(Z)$ for OR, NOT AND false.

The extrema are always calculated with $+1 > -1$.

Case 2a-2. (Convergence)

If CONN is AND or NOT OR with $TV(CONN) = F$,
OR or NOT AND with $TV(CONN) = T$,

the credibility transferred to Z is the credibility summation of the arguments of A, calculated as if each argument individually propagates credibility to Z. Credibility summation, defined here, guarantees that the sum of n summands is less than 1, provided that each summand lies between 0 and 1.

Let $0 \leq Y_i \leq 1$, then

$$\sum_{i=1}^n Y_i = Y_n + \sum_{i=1}^{n-1} Y_i - Y_n * \sum_{i=1}^{n-1} Y_i,$$

where $\sum_{i=1}^2 Y_i = Y_2 + Y_1 - Y_2 * Y_1$.

Note that credibility summation satisfies commutativity and associativity. We can now give a precise formulation of Case 2a-2, which is a case of convergence because many arguments are contributing support to one assertion.

$$ID(Z|A) = \sum_{i=1}^n [\delta(A_i) * \Delta_1(Z|A_i) * |C(A_i)|]$$

where $ID(Z|A)$ is added to $D(Z)$ if:
CONN = AND or NOT OR,
 $TV(CONN) = F$;

and $\delta(A_i) = 1$, for $C(A_i) < 0$,
 $\delta(A_i) = 0$, for $C(A_i) > 0$.

$ID(Z|A)$ is added to $I(Z)$ if: CONN = OR or NOT AND,
 $TV(CONN) = T$;

and $\delta(A_i) = 1$, for $C(A_i) > 0$,
 $\delta(A_i) = 0$, for $C(A_i) < 0$.

($\sim A_i$) substitutes for (A_i) in CONN's containing NOT. (De Morgan's theorem)

Let us sum up Case 2a in words. For a conjunction that is true, the belief that is transmitted is no more than the least belief in its arguments. For a conjunction that is false, the more arguments that are false, the more strongly supported is disbelief in the co-member of the implication. The arguments which are true make no contribution. For a disjunction, true and false are interchanged.

Case 2b. (Fanout)

$$[A \rightarrow (CONN Z_1 \dots Z_n)],$$

$$\text{or } [(CONN Z_1 \dots Z_n) \rightarrow A].$$

A is the source assertion that has changed credibility.

Z = (CONN Z₁ ... Z_n) is the target expression.
CONN is one of AND, OR, NOT AND, NOT OR.

Then, $\Delta(Z_i|A) = \Delta_1(Z_i|A) * |C(A)|$.

For AND, NOT OR,

$$I(Z) = \min \{ I(Z_i) \},$$

$$D(Z) = \max \{ D(Z_i) \}.$$

For OR, NOT AND,

$$I(Z) = \max \{ I(Z_i) \},$$

$$D(Z) = \min \{ D(Z_i) \}.$$

($\sim Z_i$) substitutes for (Z_i) in CONN's containing NOT.

In words, when a single assertion provides support to a conjunction or disjunction in an implication, first the increment or decrement to the belief in each argument must be found as if there were only the assertions A and Z_i (Case 1). Then the appropriate extrema will give the belief in the conjunction or disjunction.

Case 3. Both arguments of the implication possess multiple arguments.

Once again, we shall give rules only for conjunctions and disjunctions, reserving for Appendix III rules involving antecedents or consequents that are implications.

$$[(CONN A_1 \dots A_n) \rightarrow (CONN Z_1 \dots Z_m)],$$

$$\text{or } [(CONN Z_1 \dots Z_n) \rightarrow (CONN A_1 \dots A_n)].$$

Case 3a. If the CONN of the source expression satisfies the conditions of Case 2a-1, the source expression supports belief in the target as a single assertion, calculated as an extremum as in Case 2a-1. The resultant belief in the target expression is then calculated as a fanout from the collective source as in Case 2b.

Case 3b. In all other cases, the source expression is treated as a convergence as in Case 2a-2. Now, however, the convergence is from all A_i 's to each Z_j . After the credibilities of all the Z_j 's have been calculated, the collective credibility of Z is given by the appropriate extremum. Equivalently, this may be calculated by fanning out from each A_i to all the Z_j 's, and summing the support provided by successive A_i 's. The results are the same since credibility summation obeys the associative law.

VALIDITY, SELF-REFERENCE, SAME-SOURCE, AND EXCLUDED MIDDLE.

As S and B noted, plausible inference is only an approximation, and depends on assuming the

independence of convergent assertions just as Bayesian probability does. To the extent that the assumption is violated our calculation of credibility will be in error. This is borne out when we examine particular cases such as self-reference.

Because we are allowing arbitrary expressions of the propositional calculus, assertions of the form $[(A \vee A) \rightarrow A]$ must be considered. Care must be taken to prevent the application of credibility summation to all self-referent expressions. An expression like the above, if credibility were allowed to propagate blindly, would lead to

$$\Delta(A|A \vee A) = \Delta(A|A) + \Delta(A|A) - \Delta(A|A) * \Delta(A|A),$$

which is clearly in error if $\Delta(A|A)$ is not either 0 or 1. We therefore define $\Delta 1(A|A) = 0$ for all modes. An axiom of plausible inference is that there shall be no propagation of credibility from an assertion to itself.

A similar difficulty arises when we model sequential events by using planner-like context layers. Such modelling permits a single source of support to change credibility only slightly in a new context layer, in effect leading to adding its previous contribution to itself. This situation can be avoided in implementation by always removing the previous contribution if the present contribution is from the same source as the previous one. The contribution from other sources is then given by

$$B = (T - S) / (1 - S),$$

where B is the balance left, T is the total credibility summation, and S is the old contribution from the same source. Obviously, care must be taken not to try to remove a contribution of certainty, $S = 1$. Such an attempt is regarded as contradictory. Note that this overcomes a difficulty in S and B's monotonic system. Plausible inference is rendered non-monotonic by the use of the above equation.

The law of the excluded middle in propositional calculus may be stated as $D(A \wedge \sim A) = 1$. In plausible inference, the equation no longer holds, because $D(A \wedge \sim A) = \max\{D(A), D(\sim A)\}$. For A true, this would be simply $D(\sim A)$. This seemingly bizarre result does reduce to the normally accepted value when $|C(A)| = 1$. We could interpret this as a statement about credibility rather than truth.

[Note added after submission. Use of Shafer's equations for this case avoids the difficulty. Excluded middle falls under the case of conflicting evidence and is always false with certainty.]

DISCUSSION

The work reported here is still in development. The theoretical generalization of

propositional calculus has been completed, but applying it to realistic problems has only begun. Nevertheless, an inference engine implementing the theory has been built, and an example of trouble-shooting, exercising the system, is running. Obviously, realistic debugging is an exemplar of the general problem of pruning a large search tree. The methods suggested here for tackling the problem offer the possibility of bringing to bear all the domain specific knowledge available to the system for tree pruning, while avoiding exhaustive search.

APPENDIX I

The equivalent terms employed by Shortliffe and Buchanan are as follows:

Plausible Inference

- (1) credibility, $C(Z)$
- (2) $\Delta(Z|A)$
- (3) relevance factor, $\Delta 1(Z|A)$
- (4) $I(Z)$, $D(Z)$

MYCIN

- (1') Certainty Factor, $CF(Z,A)$
- (2') $MB(Z,A)$ or $MD(Z,A)$
- (3') $MB'(Z,A)$ or $MD'(Z,A)$
- (4') $MB(Z,A), MD(Z,A)$

Note that S and B's notation does not distinguish between an increment or decrement to an assertion and the sum of increments or decrements.

Appendix II

The case involving implication which corresponds to Case 2a is:

$$[(P \rightarrow Q) \rightarrow Z] \text{ or } [Z \rightarrow (P \rightarrow Q)],$$

where $A = (P \rightarrow Q)$ is the source assertion with changed credibility.

Then, recalling that $D(P) = I(\sim P)$,

$$I(A) = \max \{D(P), I(Q)\}, \\ D(A) = \min \{D(P), I(Q)\}.$$

This follows since:

A true is logically equivalent to $(\sim P \vee Q)$.
A false is logically equivalent to $(P \wedge \sim Q)$.

Case 2aI-1. When $TV(A) = T$, this case reduces to 2a-2, with an OR that is true. Thus

$$\Delta(Z|A) = \Delta(Z|\sim P) + \Delta(Z|Q) - \Delta(Z|\sim P) * \Delta(Z|Q),$$

where $\Delta(Z|A)$ is added to $I(Z)$.

Case 2aI-2. When $TV(A) = F$, this case reduces to 2a-1, with an AND that is true. Thus

$$\Delta(Z|A) = \Delta 1(Z|A) * [\min\{I(P), D(Q)\} - \max\{D(P), I(Q)\}],$$

where $\Delta(Z|A)$ is added to $D(Z)$.

In words, when one of the arguments of an implication is itself an implication that has changed credibility, the calculation reduces to case 2a-1 and 2a-2 by substituting the equivalent disjunction or conjunction for the implication. When A is true, the support of its arguments converges on Z, supporting its belief to the extent that NOT P and Q support belief in Z, added together as a credibility summation. When A is false, it propagates its credibility as an implication. Note, however, that when A is true, the convergence rule usually conforms to the more intuitive notion that the credibility of the implication is what is transmitted. This follows because normally the beliefs are (P true, Q true), or (P false, Q false). When that is the case, only one term contributes to the support for Z, the same result as the extremum rule.

Case 2bI. We have another fanout case where

$$[A \rightarrow (P \rightarrow Q)] \text{ or } [(P \rightarrow Q) \rightarrow A], \\ Z = (P \rightarrow Q).$$

As above, case 2bI reduces to Case 2b. For Z true, we have

$$\Delta(\sim P|A) = \Delta 1(\sim P|A) * C(A) \\ \text{added to } D(P) \text{ or } I(P) \text{ as } A \text{ is true or false.}$$

$$\Delta(Q|A) = \Delta 1(Q|A) * C(A) \\ \text{added to } I(Q) \text{ or } D(Q) \text{ as } A \text{ is true or false.}$$

$$I(Z) = \max \{ D(P), I(Q) \}. \\ D(Z) = \min \{ D(P), I(Q) \}.$$

For Z false, we have

$$\Delta(P|A) = \Delta 1(P|A) * C(A), \\ \text{added to } I(P) \text{ or } D(P) \text{ as } A \text{ is true or false.}$$

$$\Delta(\sim Q|A) = \Delta 1(\sim Q|A) * C(A), \\ \text{added to } I(\sim Q) \text{ or } D(\sim Q) \text{ as } A \text{ is true or false.}$$

$$C(Z) = \min \{ D(P), I(Q) \} - \max \{ I(P), D(Q) \}$$

APPENDIX III

Case 3aI. This is the case where either

$$[(P \rightarrow Q) \rightarrow (\text{CONN } Z_1 \dots Z_n)] \\ \text{or } [(\text{CONN } Z_1 \dots Z_n) \rightarrow (P \rightarrow Q)],$$

$$A = (P \rightarrow Q).$$

The A term reduces to Case 2aI-1 or 2aI-2, according to whether A is true or false. A then fans out to the Z terms, with or without convergence, as in case 3.

Case 3bI. This is the case for

$$[(\text{CONN } A_1 \dots A_n) \rightarrow (P \rightarrow Q)] \\ \text{or } [(P \rightarrow Q) \rightarrow (\text{CONN } A_1 \dots A_n)],$$

where $Z = (P \rightarrow Q)$ is the target assertion, and CONN is one of AND, OR, NOT AND, NOT OR. If the (CONN ...) expression reduces to a collective credibility as in Case 2a-1, 3bI reduces to 2bI.

If it is a convergence, as in 2a-2, then it is handled as in 2bI except that each term of the source is adding support to both P and Q, either to I or D as appropriate.

REFERENCES

- [1] G. Shafer, A Mathematical Theory of Evidence, Princeton University Press, Princeton, N. J., 1976.
- [2] E. H. Shortliffe and B. Buchanan, A Model of Inexact Reasoning in Medicine, Math. Biosci. 23, pp. 351-379, 1975.
- [3] E. H. Shortliffe, Computer-Based Medical Consultations: MYCIN, Elsevier, New York, 1976, Chapter 4.
- [4] J. Barnett, 'Computational Methods for a Mathematical Theory of Evidence', these proceedings.
- [5] T. Garvey, J. D. Lowrance, and M. A. Fischler, 'An Inference Technique for Integrating Knowledge from Disparate Sources', these proceedings.
- [6] G. Polya, Patterns of Plausible Inference, Princeton University Press, Princeton, N. J., 1954.
- [7] L. Zadeh, 'Fuzzy Sets', Information and Control, 8, pp. 338-353, 1965.
- [8] P. Scheffe, On Foundations of Reasoning with Uncertain Facts and Vague Concepts, Proc. AISB-80 Conference on AI, Amsterdam, July 1-4, 1980. Also Int. J. Man-Machine Studies Dec. 1980, pp. 35-62.
- [9] R. Duda, P. Hart, K. Konolige and R. Reboh '79, A Computer-Based Consultant for Mineral Exploration, Final Report, Sept. 1979, SRI International, AI Center, Menlo Park, CA. 94025.
- [10] R. Davis, Meta-Rules: Reasoning about Control and Content Reference: Reasoning about Rules, Artificial Intelligence, Vol. 15, Number 3, Dec. 1980, pp. 179-239.
- [11] L. Friedman, Trouble-Shooting by Plausible Inference, Proc. First Nat. Conf. on AI, Aug. 18-21, 1980, pp 292-294.
- [12] C. J. Rieger, The Commonsense Algorithm as a basis for Computer Models of Human Memory, Inference, Belief, and Contextual Language Comprehension, in Proc. Theoretical Issues in Natural Processing Workshop, June '75.
- [13] L. Friedman, Reasoning by Plausible Inference, Proc. 5th Conf. on Automated Deduction, Springer-Verlag, Berlin, 1980, pp. 126-142. (No. 87 in the series 'Lecture Notes in Computer Science'.)