

A Theory of Heuristic Reasoning About Uncertainty

Paul R. Cohen

Milton R. Grinberg

*Computer Science Department
Stanford University
Stanford, CA 94305*

Abstract

This article describes a theory of reasoning about uncertainty, based on a representation of states of certainty called *endorsements*. The theory of endorsements is an alternative to numerical methods for reasoning about uncertainty, such as subjective Bayesian methods (Shortliffe and Buchanan, 1975; Duda, Hart, and Nilsson, 1976) and the Shafer-Dempster theory (Shafer, 1976). The fundamental concern with numerical representations of certainty is that they hide the reasoning that produces them and thus limit one's reasoning about uncertainty. While numbers are easy to propagate over inferences, what the numbers mean is unclear. The theory of endorsements provides a richer representation of the factors that affect certainty and supports multiple strategies for dealing with uncertainty.

NOTHING IS CERTAIN. People's certainty of the past is limited by the fidelity of the devices that record it, their knowledge of the present is always incomplete, and their knowledge of the future is but speculation. Even though nothing is certain, people behave as if almost nothing is uncertain. They are adept at discounting uncertainty — making it go away. This article discusses how AI programs might be made similarly adept.

Two types of uncertainty have been studied in AI. One arises from noisy data, illustrated in speech understanding and vision programs; the other is associated with the inference rules found in many expert systems. These types of uncertainty are managed by different methods.

Noisy data are usually handled by a control structure

that exploits converging evidence. The decision to attend to a hypothesis or ignore it is related to how certain the hypothesis is: Systems are said to establish and extend "islands of certainty" (e.g., Erman, Hayes-Roth, Lesser, and Reddy, 1980). But these systems do not reason explicitly about their beliefs in their hypotheses.

When inference rules are themselves uncertain, some systems augment domain inferences with parallel *certainty inferences*. Systems such as EMYCIN (van Melle, 1980) associate *certainty factors* with the conclusions of inference rules. A rule of the form "IF A and B and C, THEN D" asserts D when A, B, and C are certain; additionally, a number may be associated with D to indicate one's belief that D follows from A, B, and C. It may be that A, B, and C, though certain, suggest but do not confirm D, in which case the number associated with D might be less than the 1.0 that usually represents certainty in such systems. If A, B, or C are uncertain, then the number associated with D is modified to account for the uncertainty of its premises. These numbers are given different names by different authors; we refer to them as *degrees of belief* (Shafer, 1976). The functions that propagate degrees of belief over inferences are called *combining functions*. Domain rules are assigned a priori degrees of belief and the purpose of the combining functions is to faithfully represent the intent of each in the context in which it is eventually used. Some systems propagate not one degree of belief, but two, indicating a range of certainty. In

all cases, one's certainty in a hypothesis is represented only by a numerical degree of belief.

Problems with Current Approaches to Uncertainty

There are serious limitations to current numerical approaches to reasoning under uncertainty. Some approaches use just a single number to represent a degree of belief, but, as Quinlan (1982) points out, "the single value tells us nothing about its precision." Another problem with single numbers is that they combine evidence for and against a proposition, and so one cannot distinguish between disbelief and a lack of evidence pro or con (Shafer, 1976). Various schemes have been used to overcome these representational deficits, such as ranges instead of point values and separate measures of belief and disbelief (see Quinlan, 1982, for a review.)

Most systems use a variant or generalization of Bayes's theorem to derive the degree of belief in a conclusion from the degrees of belief of the preconditions. Unfortunately, Bayes's theorem requires masses of statistical data in addition to the degrees of belief in preconditions. Almost always, subjective expert judgments are used in place of these data, with the risks of inaccuracy and inconsistency (see Shortliffe and Buchanan, 1975; Duda, Hart, and Nilsson, 1976; Shortliffe, Buchanan, and Feigenbaum, 1979, discuss the reasons for success and failure in several medical programs that use Bayes's theorem.)

These are well-known, documented problems with numerical approaches to reasoning about uncertainty. The remainder of this section discusses the representational problems in more detail. In particular, it proposes that numerical approaches to reasoning under uncertainty are restricted because the set of numbers is not a sufficiently rich representation to support considerable heuristic knowledge about uncertainty and evidence.

Numerical degrees of belief in current AI systems are of two kinds: those specified *initially* as qualifications of domain

inference rules and those that are *derived* from the initial numbers as the system reasons. Initial numbers are usually supplied by domain experts; for example, an expert investment counselor may associate a *degree of belief* of 0.6 with the inference that advanced age implies low risk tolerance. It is not always clear what the 0.6 means. It may mean that 60% of the distribution of the elderly people in a sample have low risk tolerance. More often, the 0.6 represents an expert's *degree of belief* that a person has low risk tolerance if he or she is old. The number is presumably a summary of the reasons for believing and disbelieving the inference; but once summarized, these reasons are inaccessible. This is one reason that explanations in expert systems are limited to a recitation of the inferences that led to a conclusion. Current systems explain how they arrived at a conclusion, but not why they tend to believe (or disbelieve) it. None of these systems appears able to interpret its degrees of belief.

The second kind of numerical degrees of belief that are found in AI systems are those derived by reasoning. A general schematic of the derivation of degrees of belief is shown in Figure 1. At the top of the figure are two domain inferences from investment counseling: Rule 1 states that advanced age implies low risk tolerance, and rule 2 infers that the client needs a high proportion of bonds if he (or she) has low risk tolerance. Associated with each rule is an initial degree of belief (0.6 and 0.8, respectively) supplied by the domain expert. These numbers are combined with the degrees of belief of their rule's premises to produce a *derived* degree of belief in the conclusion. For example, if it is *certain* that risk tolerance is low, then 0.8 represents the undiluted degree of belief in the conclusion that the client should have a preponderance of bonds. But if—as in Figure 1—it is less than certain that risk tolerance is low, then the degree of belief in the conclusion about bonds should be less than 0.8. In Figure 1, this premise is not fully believed and the certainty in the conclusion is represented by the product of 0.8 and 0.6. The function that combines these two numbers—in this case by multiplication—is called a combining function.

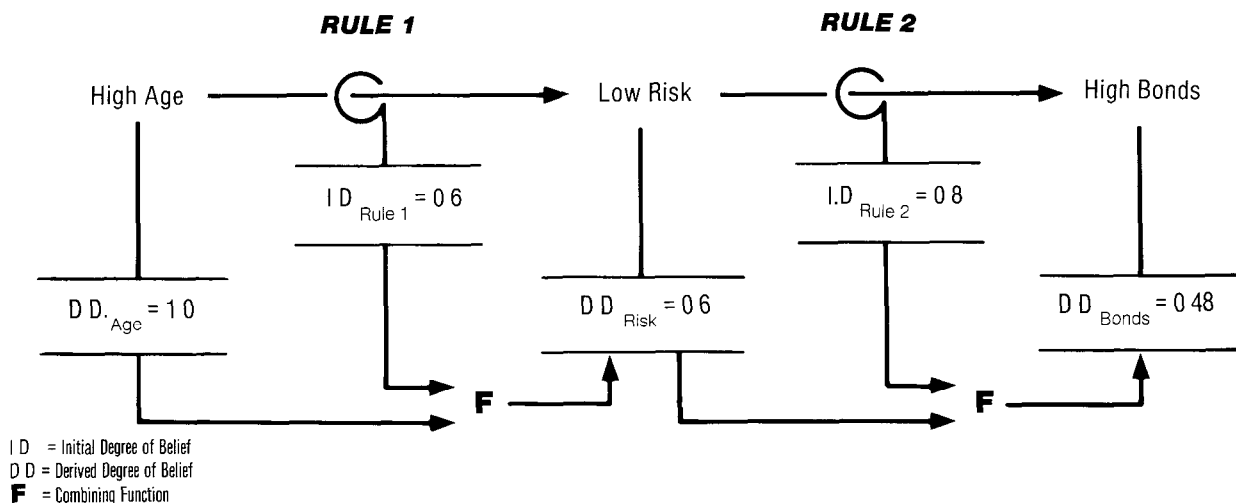


Figure 1 The derivation of degrees of belief

Initial degrees of belief—those associated with inference rules by the system builder—are numerical constants, specified before the system is used. In contrast, the derived degrees of belief that modify individual propositions are variables. In a sense, inferences are made in the *context* of the degrees of belief in their preconditions, and combining functions make the initial degrees of belief, associated with each inference, sensitive to these contexts (see Fig. 2). We use the word *context* to convey the idea that situations arise dynamically that affect one's certainty. Later we propose that the definition of context must be broadened to include more than just the degrees of belief of preconditions

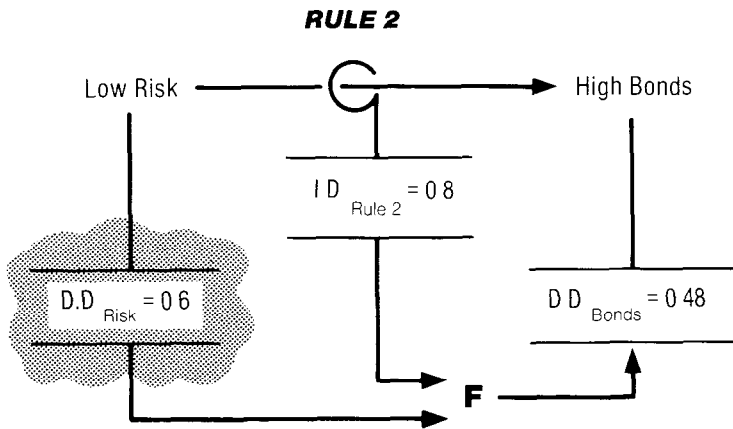


Figure 2 D D Risk is the context in which rule 2 is invoked
F makes the rule sensitive to its context

Most AI systems that reason about uncertainty follow the approach of Figure 1 to derive degrees of belief from initial degrees of belief and uncertain antecedents. In these systems the scheme is adhered to with little or no variation. All evidence is processed in exactly the same way: It propagates through a combining function to the conclusion it supports. Current systems are unable to treat different kinds of evidence differently. They do not take some evidence "with a grain of salt" and other evidence with judicial solemnity, and the reason is obvious: Since evidence is nothing more than a proposition with an associated number, there is no way to tell whether it warrants scrutiny other than by examining the number. This permits only a limited ability to discriminate kinds of evidence. For example, EMYCIN discriminates two kinds of evidence: If the evidence for a conclusion does not have a certainty factor (CF) of at least 0.2, the conclusion is not asserted; otherwise, it is asserted and qualified by a derived CF. This is the extent of the system's reflection on its evidence. But what more can be expected when its only information about its evidence is a number? If there are several kinds of evidence, then the numerical representation is inadequate; for example, one cannot expect to discriminate eyewitness evidence from circumstantial, hearsay, or photographic evidence on the basis of their numerical degrees of belief.

If the degree of belief associated with a proposition does not represent the kind of evidence that supports the proposition, a system will be unable to reason about different types of evidence differently. An extension of this problem is that the same evidence may be treated differently in different contexts. For example, eyewitness evidence is usually trustworthy and is often the basis for criminal convictions. But recently, a man was convicted on the basis of eyewitness testimony of a crime committed by another man of very similar appearance. The testimony against the first man was compelling, but it lost its force when the second man was apprehended. In one context, before the second man was found, the eyewitness evidence outweighed the first man's alibi; in another context, the weights of the two kinds of evidence were reversed. It is difficult to imagine how current AI systems could be made to reason this way. Eyewitness testimony relies on some assumptions, notably that individuals have distinctive appearances. If this assumption is violated, as it was in the example, then the evidence loses force. Since numerical degrees of belief do not represent these assumptions, they are insufficient for the task of changing the weight of evidence should an assumption be violated

Numerical degrees of belief are adequate if one intends only to propagate them over inferences without reflection. Indeed, they do not support much reflection. But the example illustrates that it is often necessary to reason *about* evidence¹ and the contexts in which it is used. Two problems must be solved to achieve this. The first is to represent explicitly many aspects of evidence besides "how much" it is believed. The other problem is to develop additional methods for reasoning with this richer information. Unlike combining functions, which automatically treat all evidence in exactly the same way, heuristic reasoning about uncertainty is contingent on aspects of evidence and context.

The earlier legal example raised the issue of justifying, instead of just quantifying, one's uncertainty. This emphasis leads to AI research on reason maintenance. Doyle (1979) has developed a *truth maintenance system* (TMS), a theory of how a program might record its reasons for believing a hypothesis. Reasons are used to construct explanations and to revise sets of belief. Every inferred node *n* is stored with its *justifications*—roughly, a record of the nodes on which *n* depends. However, Doyle makes little differentiation between kinds of justifications: "Although natural arguments may use a wealth of types of argument steps or justifications, the TMS forces one to fit all these into a common mold" (p. 239). Doyle is primarily interested in whether a node has support, not in what kind of support it has. This is equivalent to saying, in the legal case, that there are reasons for a conviction without saying what the reasons are. Nevertheless, the idea of a reason maintenance system can be adapted to the task of recording and reasoning about kinds of evidence and their reliability and diagnosticity

¹Hence, the title of this article.

An Example of Intelligent Reasoning About Uncertainty

The theory of endorsements is developed in the context of an example taken from the domain of portfolio management. An expert system called FOLIO (Cohen and Lieberman, 1983) interviews a client to determine his (or her) asset structure and makes recommendations about how to allocate the client's assets. In the example here, FOLIO has already calculated its estimate of the client's pretax interest income, **Estimated-pti**. This parameter is derived from the client's description of his investments and from an assumption about the interest rate on fixed-income investments. It should be consistent with the client's own estimate of his pretax interest income, **Reported-pti**. Figure 3 illustrates a fragment of FOLIO's reasoning in the absence of uncertainty: **Estimated-pti** corroborates **Reported-pti** and produces **Resolved-pti**; this, in turn, is one of the components of **Adjusted-Gross-Income**, from which Tax-bracket can be estimated. However, even this small fragment provides a vehicle for much reasoning about uncertainty.

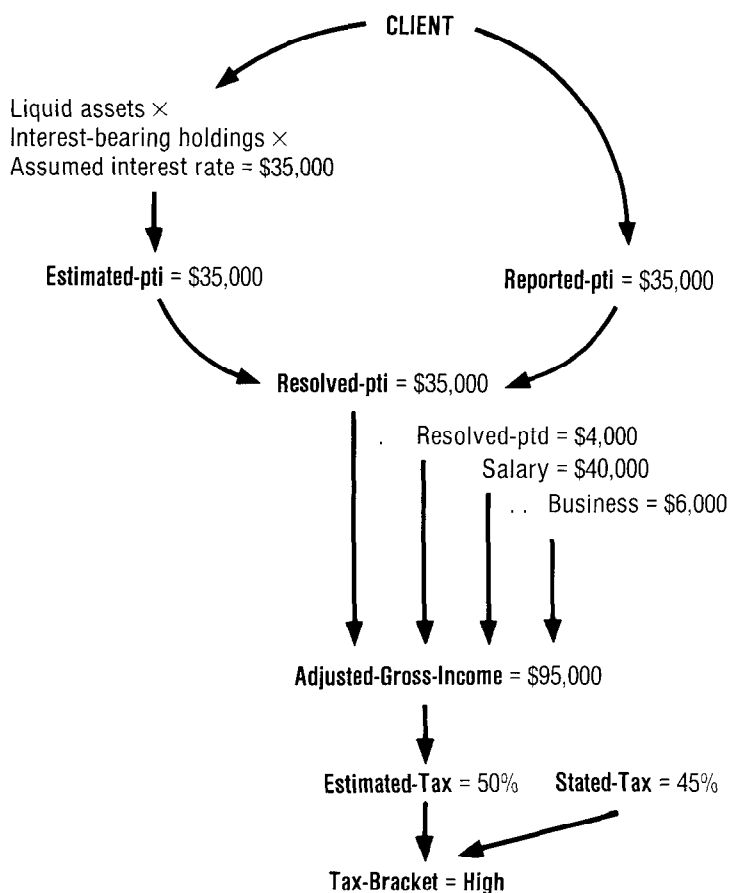


Figure 3 Derivation of Tax Bracket

Imagine the several lines of reasoning that are possible if **Estimated-pti** does not corroborate **Reported-pti** but differs by \$5,000:

- *I estimated the client's pretax interest income, **Estimated-pti**, to be \$35,000. The client said that **Reported-pti** = \$40,000, so there's a \$5,000 discrepancy. Is it important? If not, one approach to the uncertainty is to resolve **Resolved-pti** to be either \$35,000 or \$40,000, or maybe split the difference between them. The prerequisite for this approach is to know whether \$5,000 is a big enough chunk of \$35,000 to warrant a more exact resolution. (For some purposes, e.g., computing income tax, \$5,000 is too much to ignore. Other goals, such as estimating whether the client is producing more income than necessary, would probably not be impeded.)*
- *Another approach is to discard one of the two estimates. I recall that **Estimated-pti** was based on an assumption about interest rates, but I used today's rates for my estimate, and the rates jump around all the time, so maybe I should prefer the client's estimate to my own.*
- *Alternatively, I could admit that I don't know which of **Reported-pti** and **Estimated-pti** is preferable. In that case, I could propagate them both.*

Let us assume a decision is made to accept **Reported-pti** as the value for **Resolved-pti** based on the arguments above. Then reasoning about the value of **Adjusted-Gross-Income**, and ultimately about **Tax-bracket**, might continue as follows:

- *The next thing I need to know is the client's adjusted gross income, but one of its components is pretax interest income, **Resolved-pti**, and I'm uncertain about it. What does this imply about my certainty of **Adjusted-Gross-Income**? There was a conflict between **Reported-pti** and **Estimated-pti** that I resolved by choosing the former. Am I still satisfied with that resolution? Is there a better resolution? Yes. **Adjusted-Gross-Income** is derived from many other parameters, of which **Resolved-pti** is but one, so the \$5,000 discrepancy is not so large in this new context, and I can discount my uncertainty about **Resolved-pti**, at least for the purposes of this computation. In this context, I prefer to discount the uncertainty by averaging the values of **Estimated-pti** and **Reported-pti** rather than by rejecting one of them outright.*
- *Now, given a value for **Adjusted-Gross-Income**, I can compute a tax bracket. I estimate the client's bracket to be 50%, which conflicts with the client's estimate of 45%. Is the difference important? Can I trade my uncertainty about the exact tax bracket for certainty about a less precise hypothesis? Yes, I am prepared to say that the tax bracket is not known with certainty, but it is certainly "high."*

An Alternative Approach

Three central themes in the theory of reasoning about uncertainty are illustrated in this example. First, repeated attempts are made to discount uncertainty. For example,

the discrepancy between **Reported-pti** and **Estimated-pti** might be explained by the fact that the latter was based on an assumption about interest rates. In general, a system that reasons about uncertainty should be expected to try to explain it away. The second theme is that there are many methods for discounting uncertainty. The initial *assumption* about interest rates was preferable to uncertainty about the rate of return on each of the client's investments. Assumption is a method for discounting uncertainty, and it is an instance of a more general method that involves *trading specificity for certainty*. This general approach was illustrated again when the average of **Estimated-pti** and **Reported-pti** was used to estimate **Adjusted-Gross-Income**. In this case, knowledge of *which* of **Estimated-pti** and **Reported-pti** is more accurate is traded for certainty that their average is accurate, or at least accurate enough for the purpose of computing **Adjusted-Gross-Income**. This example also illustrates a second very general method, namely, accepting a hypothesis as *sufficiently certain* for a particular task. It should be possible to know enough about the value of **Resolved-pti** to confidently compute **Adjusted-Gross-Income** without implying that **Resolved-pti** is known with certainty. The third theme is that reasoning about uncertainty is knowledge intensive. Besides the very general knowledge used to discount uncertainty, there is a vast amount of very specialized domain-specific knowledge that needs to be applied. For example, it was the knowledge that interest rates fluctuate that was the basis for rejecting **Estimated-pti** as an estimate of **Resolved-pti**.

A representation of states of certainty that is rich enough to support the inferences needed to discount uncertainty is required. This representation must make explicit the knowledge that is only summarized in numerical degrees of belief. For example, if the estimates based on interest rates are unreliable, this should be explicitly stated. There is no advantage and considerable disadvantage to summarizing the mistrust of this kind of estimate in a degree of belief of, say, 0.4. Different kinds of evidence should be distinguished by an explicit record of what makes them different, not by numbers between 0.0 and 1.0.

Endorsement and Endorsers

The explicit marking of factors relating to one's certainty is called an *endorsement*. The process is best introduced by an analogy. A piece of work in a bureaucracy proceeds from one stage to the next contingent on the endorsement of a bureaucrat (see Fig. 4). The job must satisfy certain (typically formal) requirements before it is endorsed at any given level. Imagine, in place of bureaucrats watching over a job, a collection of rules watching over the development of a line of reasoning. Each endorses a step in the argument if it satisfies certain requirements. Whenever a domain rule is used, its conclusion accrues one or more endorsements. Thus,

endorsements are just records that a particular kind of inference has taken place, and *endorsers* are just the computations that assert the records. Bureaucrats can require a job to be cleared by lesser bureaucrats before they even consider it; for example, a city council won't consider a proposal unless it is cleared by the planning department, and they won't look at it unless the lawyers have approved it first. Similarly, an endorser may require the conditions of a rule to have a certain level of endorsement before it will endorse the conclusion of the rule. For example, one endorser might endorse the conclusions of a rule only if the conditions were themselves endorsed as parameters derived from rules that do not introduce uncertainty, such as simple arithmetic transformations. Most conclusions accrue several, more or less stringent endorsements. The certainty of a hypothesis can be represented as its strongest endorsement. In terms of the bureaucracy analogy, one's confidence in a job is proportional to the degree of scrutiny and level through which it has passed.




Figure 4 Bureaucrats endorse proposals and discard those with insufficient endorsements

Endorsement is similar to recording justifications in a truth maintenance system (Doyle, 1979), but with a crucial difference. In the TMS, a justification is used to decide whether a conclusion has support, but the kind of support is irrelevant (see Fig. 5a). Endorsements, however, record aspects of inferences that are relevant to reasoning about their certainty. There are many different kinds of endorsements, corresponding to evidence for and against a proposition, and one's confidence in a conclusion depends on its *pedigree*— its endorsements and those of its preconditions. Implicit in this is the idea that endorsements can be ranked. Preference of the endorsement of one hypothesis over the endorsement of another is equivalent to having more confidence in the one hypothesis over the other. For example, in most

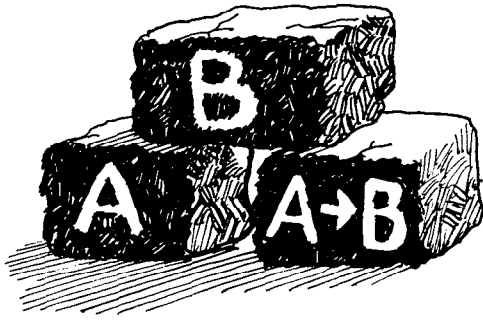


Figure 5a Truth-maintenance-style support for Proposition B

contexts, eyewitness testimony is preferable to circumstantial evidence, direct evidence to indirect, corroboration to contradiction, and inference to assumption. What is meant by one kind of evidence being preferable to another is that a conclusion endorsed as having the one kind of evidence for its support is more certain than a conclusion endorsed as supported by less preferred evidence (see Fig. 5b) Much knowledge is needed to rank endorsements. For example, it is difficult to know whether the eyewitness account of a drunk is more certain than a dozen “respectable” but circumstantial anecdotes, but the ability even to pose the question (if not answer it) is evidence for the role of world knowledge in weighing evidence.




Figure 5b Endorsement support for Proposition B

Just as degrees of belief are propagated over inferences by combining functions, so must endorsements be propagated over inferences by heuristics. Given that one proposition implies another, and the premise has a set of endorsements, what endorsements should the conclusion accrue (see Fig. 6)? A set of rules is needed to propagate endorsements over inferences. These serve the same purpose as combining functions, to make a rule sensitive to the context in which it is used, but the context is now a set of endorsements instead of a single degree of belief. A default rule is that the endorsements of the premise should all propagate to the conclusion. In fact, there are interactions between endorsements that complicate this picture. For example, if a parameter is derived by taking a *central value* of several others (e.g.,

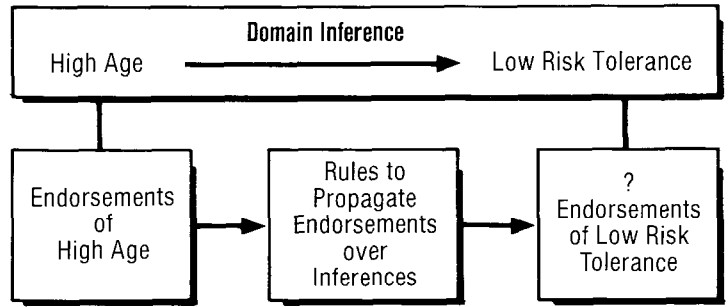


Figure 6

an arithmetic mean) and if one of the component values is thought to be an **extreme** value, the endorsement **extreme** usually would *not* apply to the central value. Each domain of expertise is expected to have numerous idiosyncratic rules for propagating endorsements over inferences.

Endorsements can also be used to resolve uncertain values and discount uncertainty, although this, too, requires masses of general and special knowledge. A general method for resolving uncertain values is to choose a central or representative value in their stead. Alternatively, the one with the best endorsement might be selected; for example, the “credentials” of disputants can help decide between their claims—until recently, the word of a congressman outweighed that of a convicted swindler. Other methods are specific to the kinds of uncertainty they resolve. For example, if you can’t resolve whether to withdraw \$50 or \$100 for an evening’s entertainment, it’s best to get \$100. A method of intermediate generality applies to summary parameters like **Adjusted-Gross-Income**: If a numerical parameter is a sum of several other *independent* parameters, then uncertainty about any one of the component parameters can be ignored for the purposes of calculating the summary value.

Resolution of an uncertainty does not imply that the resolution is certain; it can be *certain enough* for one purpose but not for another. For example, an average of **Estimated-pti** and **Reported-pti** may be a sure enough estimate of pretax interest for the purpose of computing **Adjusted-Gross-Income** but not for deciding whether a portfolio generates too much interest income. The idea of *complete certainty* is an artifact of numerical representations of degree of belief. It is more productive to think in terms of *certainty with respect to a task*; in fact, it makes little sense to speak of certainty except with respect to a task.

The meaning of endorsements is determined by how a system reasons with them. They have an operational semantics. Words like *direct* or *corroboration* are just the names of endorsements, chosen because their English meaning corresponds to some characteristic of evidence. One part of the meaning of the **corroboration** endorsement is the interaction of evidence to which the word is applied: A proposition has the support of two or more pieces of evidence. A second component of the meaning of the **corroboration** endorsement is how propositions that carry it are ranked relative to other propositions; for example, a proposition endorsed

with **corroboration** is usually or always preferred to one endorsed with contradiction. A third component of the meaning of the **corroboration** endorsement is how it interacts with other endorsements. For example, if just *one* of two corroborating pieces of evidence is endorsed by **inaccurate**, this endorsement need not apply to the proposition it supports; but if all corroborating evidence is endorsed by **inaccurate**, one might endorse the proposition they support with some warning of the possibility of compounding errors.

We expect that each domain of expertise has a characteristic set of endorsements and a set of methods that define what they mean. These methods include ranking rules, rules for propagating endorsements over inferences, and rules for discounting uncertainty. These combine to reason about nodes or conclusions in an inference net (such as that shown in Fig. 3). An uncertain conclusion can be resolved in one of four ways:

1. A node's endorsements can be judged to be sufficient for the goal at hand. This is not equivalent to saying that the value of the node is certain, only that it is not uncertain enough to warrant other action.
2. An endorsement of an earlier node, although sufficient for some earlier goal, is judged to be insufficient for the current goal. The earlier value is retrieved and reconsidered in the current context, and a new endorsement (or possibly the old one) is assigned. This involves backtracking.
3. There is uncertainty about the value of the current node, but it is discounted by (a) picking the value that has the highest endorsement or (b.) generating a new value according to some heuristic method
4. The uncertainty of the current node cannot be resolved in a way that preserves a minimum endorsement, so the multiple values of the current node are propagated on to the next node.

An Example of Reasoning with Endorsers

How endorsements facilitate the reasoning in the example is now illustrated. Figure 3 (presented earlier) shows the inferences that would be made if no uncertainty arose. In the case that **Estimated-pti** and **Reported-pti** differ, the following reasoning takes place:

- *I estimated the client's pretax income, **Estimated-pti**, to be \$35,000. The client said that **Reported-pti** = \$40,000. Is it important? Not really. I can resolve **Resolved-pti** to be either \$35,000 or \$40,000 or maybe split the difference between them.* This corresponds to case 3.b above and the use of the general heuristic that if a conflict is not important, resolve it arbitrarily. The resulting value would be endorsed as arbitrarily derived and presumably would not rank as highly as better justified resolutions
- *Another approach is to discard one of the two estimates. **Estimated-pti** was based on an assumption about in-*

*terest rates, and these jump around daily, so maybe I should prefer **Reported-pti**. This corresponds to case 3.a above. There are actually two justifications here for discarding **Estimated-pti**, or, rather, it carries two damning endorsements that can be used to abandon it. First, its own endorsement includes an **Assumption**, while **Reported-pti**'s endorsement is **Direct evidence**, and the latter is preferred. Second, its endorsement records that **Estimated-pti** was derived from information with **low predictive power**—today's interest rates. Both of these negative endorsements are cause to prefer **Reported-pti** and endorse it with **Resolve-by-endorsement**.*

- *Alternatively, I could admit that I don't know which of **Reported-pti** and **Estimated-pti** is preferable. In that case, I could propagate both.* This corresponds to case 4 above. It is an expensive strategy, since it may eventually result in a combinatorial explosion like that found in decision analysis trees (Schlaifer, 1969). In this case, the alternatives maintain their endorsements, and the hope is that they can be resolved in the context of information available at a node farther down the inference net. For example, the uncertainty that arises from the conflicting parameters may not be resolvable at the node **Resolved-pti** but can be resolved at **Adjusted-Gross-Income** by the *summary parameter heuristic* that was mentioned above: Uncertainty about the value of one of many independent components of a summary value can be discounted. (This kind of reasoning corresponds to case 3 b above.)
- *The next thing I need to know is the client's adjusted gross income, but one of its components is pretax interest income, **Resolved-pti**, and I'm uncertain about it. What does that imply about my certainty of **Adjusted-Gross-Income**? There was a conflict between **Estimated-pti** and **Reported-pti** that I resolved by choosing the former. Am I still satisfied with that resolution? Is there a better resolution? Yes. Since **Adjusted-Gross-Income** is derived from many other parameters, of which **Resolved-pti** is but one, I am prepared to discount my uncertainty about **Resolved-pti**, at least for the purposes of this computation. This corresponds to case 2 above. At the point of computing **Adjusted-Gross-Income**, the endorsement of **Resolved-pti**, which is **Resolved-by-endorsement**, is unsatisfactory, so backtracking occurs to find the original conflict between **Estimated-pti** and **Reported-pti** and consider the conflict in the context of computing **Adjusted-Gross-Income**. As it happens, the uncertainty can be resolved by the summary-parameter heuristic, and this results in an endorsement that is preferable to **Resolved-by-endorsement**.*
- *Now, given a value for **Adjusted-Gross-Income**, I can compute a tax bracket. I estimate the client's bracket to be 50%, which conflicts with the client's estimate of 45%. . . I am prepared to say that the bracket is not known with certainty, but it is certainly "high."* This is another example of case 3 b. Note that the method is very general—to discount small differences—but it cannot be implemented with confidence unless there

is some knowledge that the difference is actually small. In this case, it must be known that 5% is not a big discrepancy between tax estimates. But this cannot be so assured unless the purpose of the estimate is known. Certainly, the IRS doesn't regard 5% as a small difference! Thus, the endorsement for this resolution (which might be called **Discount discrepancy**) is suspicious. It is a resolution of uncertainty but one that indicates to later processes that they might want to reconsider the original estimates, as in case 2 above

Summary

A heuristic approach to reasoning about uncertainty and a representation that supports it has been presented. Reasoning about uncertainty is a knowledge-intensive task, and domain experts have their own heuristics for dealing with typical kinds of uncertainty. Endorsements are records of information that affect one's certainty. This includes the kind of evidence that is available (direct, corroborative, eyewitness, etc.) and the kinds of methods used to produce the current hypothesis from uncertain predecessors (averaging, eliminating one of the conflicting estimates, etc.). Endorsements can be ranked, and a hypothesis with a superior endorsement is more certain than another that is less well endorsed. Endorsements can also be propagated over inferences, but in a manner that is sensitive to the context of the inference. Thus, endorsements are an alternative representation to numerical degrees of belief, one that supports sophisticated reasoning about uncertainty.

References

- Cohen, P. R., and Lieberman, M. D. 1983. A report on FOLIO: An expert assistant for portfolio managers. Unpublished report. Computer Science Department, Stanford University.
- Doyle, J. 1979. A truth maintenance system. *Artificial Intelligence* 12: 231-272.
- Duda, R. O., Hart, P. E., and Nilsson, N. 1976. Subjective Bayesian methods for rule-based inference systems. Technical Note 124, Artificial Intelligence Center, SRI International, Menlo Park, CA.
- Erman, L., Hayes-Roth, F., Lesser, V., and Reddy, D. R. 1980. The HEARSAY-II speech understanding system: Integrating knowledge to resolve uncertainty. *Computing Surveys* 12(2):213-253.
- Quinlan, J. R. 1982. INFERNO: A cautious approach to uncertain inference. RAND Note N-1898-RC, RAND Corporation, Santa Monica, CA.
- Schlaifer, R. 1969. Analysis of decisions under uncertainty. New York: McGraw-Hill.
- Shafer, G. 1976. A mathematical theory of evidence. Princeton, NJ.: Princeton University Press.
- Shortliffe, E. H., and Buchanan, B. G. 1975. A model of inexact reasoning in medicine. *Mathematical Biosciences* 23:351-379.
- Shortliffe, E. H., Buchanan, B. G., and Feigenbaum, E. A. 1979. Knowledge engineering for medical decision making: A review of computer-based clinical decision aids. *Proceedings of the IEEE*, 67(9):1207-1224.
- Van Melle, W. 1980. A domain-independent system that aids in constructing knowledge-based consultation programs. Report STAN-CS-80-820, Department of Computer Science, Stanford University.

Advertisement

JPL Artificial Intelligence Research

The Jet Propulsion Laboratory has positions available in Artificial Intelligence Research. These positions require a PhD in either Computer Science or Artificial Intelligence with at least two years experience. Candidates must have backgrounds in AI planning, problem-solving, intelligent reasoning systems or expert systems and demonstrate the capacity to implement an AI system.

Please submit resume to Professional Staffing, Department J35.

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
Equal Opportunity Employer M/F