# BELIEF MAINTENANCE: AN INTEGRATED APPROACH TO UNCERTAINTY MANAGEMENT

Kathryn B. Laskey and Paul E. Lehner

Decision Science Consortium, Inc.
1895 Preston White Drive
Reston, Virginia 22091

George Mason University
4400 University Drive
Fairfax, Virginia 22030

**Abstract**

Belief maintenance represents a unified approach to assumption-based and numerical uncertainty management. A formal equivalence is demonstrated between Shafer-Dempster belief theory and assumption-based truth maintenance extended to incorporate a probability calculus on assumptions. Belief propagation through truth maintenance automatically and correctly accounts for non-independencies among propositions due to shared antecedents. Belief maintenance also incorporates an ability to represent and reason with defaults. The result is a framework for non-monotonic reasoning about the application of a quantitative uncertainty calculus.

## 1.0 INTRODUCTION

Automated reasoning systems must operate with incomplete knowledge of the state of the world. Much of the work of problem solving or inference lies in structuring exploration of the system's world to reduce this uncertainty. Two general approaches to uncertainty management have become popular. These approaches--symbolic truth maintenance and numeric belief propagation--have been portrayed as rivals in a sometimes acrimonious debate. Yet each has an important role to play in a comprehensive inference strategy.

Symbolic truth maintenance [Doyle, 1979; deKleer, 1986] is based on the idea of extending ordinary truth-functional logic to allow the incorporation of defaults or assumptions. A system based on such a logic can set default values for uncertain propositions and reason as if these values were known, but revise its defaults when they give rise to a contradiction. This capability mimics the non-monotonic character of intelligent human reasoning. Truth maintenance increases search efficiency by permitting a control strategy that minimizes regeneration of previously considered search paths. Uncertainty is represented entirely qualitatively: what is known about propositions is whether they are proven true, assumed true, unknown, assumed false, or proven false in a given context.

In the quantitative approach, uncertainty about a proposition's truth value is expressed by a number or numbers, typically in the range between 0 and 1. Inference rules also have numbers associated with them, indicating how strongly belief in the premises warrants belief in the conclusions. Various calculi have been proposed for propagating beliefs in chains of inference.

We argue that a significant advantage may be had by combining symbolic and numeric uncertainty management. Specific advantages include the following.

*Computational aspects of numeric reasoning*. The numeric approach has been criticized on efficiency grounds. The number of possible combinations of truth values grows exponentially with the number of propositions in a system. A numeric calculus maintains propositions, no matter how improbable, unless they are explicitly proven false. Of course, sophisticated computational architectures and control strategies can reduce inefficiency (e.g., by not exploring consequents of improbable propositions). But additional control strategies open up with the possibility of applying non-monotonic qualitative reasoning to the *application* of an uncertainty calculus. For example, the control strategy might make an *assumption*, which could be later retracted, assigning a truth value of false to an improbable proposition. Or the structure of an inference network could be subjected to qualitative reasoning by making retractable conditional independence assumptions.

*Control of reasoning*. An important role of truth maintenance is to control reasoning--to decide what to do next. Additional possibilities for control are opened up by including a numeric uncertainty calculus. An example cited above is to avoid using inference rules with improbable antecedents. Another strategy involves searching for information that will likely distinguish between two uncertain but competing hypotheses.

*Conflict resolution*. The two uncertainty management traditions have very different approaches to dealing with conflict [Cohen, Laskey and Ulvila, 1987]. In the symbolic tradition, conflicting conclusions indicate that one of the lines of reasoning is faulty.

Defaults must be changed to restore consistency. By contrast, probabilistic and other numerical systems regard conflict as an inevitable consequence of an imperfect correlation or causal link between evidence and conclusion. Probabilistic reasoning regards inference as weighing a balance of positive and negative evidence; symbolic reasoning adjusts the set of accepted defaults so that no simultaneous arguments exist for both a proposition and its complement. Intelligent reasoning, we feel, combines aspects of both these viewpoints. We must be able to entertain conflicting lines of argument, adjudicating between them based on the strength of each. On the other hand, when the conflict becomes too great, we begin to suspect a problem with one or the other argument, and examine our implicit beliefs for possible revision.

This paper represents a preliminary effort toward unifying symbolic and numeric reasoning. It has been implemented, but only on small-scale problems. Much work remains in developing control strategies and algorithms to circumvent the computational complexity associated with large belief networks. But the structure described here represents an important step toward understanding the relationship between probabilistic and belief function models, truth maintenance, and non-monotonic logics.

## 2.0 THE ATMS AND SHAFER-DEMPSTER BELIEFS

This section introduces an approach to uncertainty management that unifies the symbolic and numeric approaches. In particular, it is demonstrated that extending assumption-based truth maintenance [deKleer, 1986] to allow a probability calculus on assumptions leads to a belief calculus that is formally equivalent to Shafer-Dempster theory. By appropriate construction of inference rules, probabilistic models emerge as a special case. An advantage of the framework presented here is that nonindependencies are automatically and correctly accounted for in the belief computations.

Let us introduce the theory with an example used by Shafer [1987] to illustrate belief function theory. Suppose we receive a report $r$ that a proposition $q$ is the case. There is some probability, say .6, that the source is reporting reliably; otherwise, with probability .4, the report bears no relation to the truth of $q$. In the first case, the report implies $q$; in the second case, both $q$ and its negation are consistent with the report. For Shafer, belief in a proposition is defined as the probability that the evidence implies its truth. Thus, the above evidence justifies a .6 degree of belief in the proposition $q$.

Another way of looking at this example is in terms of an inference rule ($r \rightarrow q$) that may or may not be valid. A symbolic default reasoning system might incorporate a default assumption that the rule is operating correctly (i.e., the source is reliable). Receipt of evidence against $q$ would necessitate dropping either this default or one of the assumptions underlying the conclusion of $\neg q$. The belief calculus, on the other hand, simultaneously maintains belief that the rule is and is not valid, using a number between zero and 1 to encode the strength of belief in rule validity.

Simultaneously maintaining belief in inconsistent propositions is impossible in most default reasoning systems. When beliefs become inconsistent, these systems must explicitly change some of their defaults to regain consistency. Combining beliefs and defaults is therefore difficult with traditional truth maintenance. But the ATMS [deKleer, 1986] is explicitly designed to reason with multiple contexts. Each proposition is tagged with a *label* that represents the contexts in which it can be proven. Contexts are represented by special symbols which deKleer calls assumptions. We depart from deKleer's terminology and use the term *tokens* to refer to these special symbols, because we prefer the word assumption to retain the connotation of a proposition that, although unproven, has been declared to be in the set of believed propositions. The job of the ATMS is to maintain a parsimonious representation of each proposition's dependence, either directly or through chains of inference, on tokens. Like deKleer, we define an *environment* to be a set of tokens and a *context* to be the entire set of propositions derivable in a particular environment. A proposition's label contains a list of environments in which it can be derived. There may be proofs for a proposition under several mutually inconsistent environments; similarly, proofs for a proposition and its negation in different environments may be entertained simultaneously.

Consider again the inference from report $r$ to the truth of the reported proposition $q$. In the ATMS, the "noisy" inference rule may be encoded as follows. The token V is introduced to represent rule validity, i.e., the proposition that the source is reporting reliably. (Following deKleer, we use uppercase letters to represent tokens and lowercase letters to represent other propositions). The original noisy rule is replaced by the rule $r \wedge V \rightarrow q$. When this rule fires, the token V is added to the label of the ATMS node associated with $q$, indicating that $q$ is valid in any context in which V is true.

If this token is treated as a default assumption, then the proposition $q$ is believed until this default becomes inconsistent with the set of known propositions. Alternatively, V may be assigned a probability, interpreted as the probability that $q$ can be proven--that is, its Shafer-Dempster belief.

Thus, belief maintenance is based on a simple principle: if probabilities are assigned to tokens, these imply probabilities on the labels for propositions. The probability of a label can be interpreted as the probability that the associated proposition can be proven, and is equivalent to its Shafer-Dempster belief.

## 3.0 A BELIEF MAINTENANCE SYSTEM

Our *belief maintenance system* combines an ATMS with a module for computing the probabilities of ATMS labels, or, equivalently, the Shafer-Dempster beliefs of the associated propositions. Belief maintenance is capable of representing the full generality of the Shafer-Dempster calculus. The ATMS automatically keeps account, in symbolic form, of the propagation of beliefs through chains of inference, nonindependencies created through shared premises, and inconsistent combinations of tokens. The belief computation module incorporates all this information to compute correct Shafer-Dempster beliefs when requested. Adding to this framework the capability to represent and compute beliefs with defaults results in a fully integrated symbolic and numeric uncertainty management framework.

Readers familiar with the basics of belief function theory and assumption-based truth maintenance will more easily follow the following presentation. Laskey and Lehner [1987] include a concise introduction to both theories [see also deKleer, 1986; Shafer, 1976].

### 3.1 Combining Beliefs and Chaining Inference Rules

We have shown how a single uncertain inference rule could give rise to a Shafer-Dempster belief. But interesting inference problems involve many propositions, linked together by complex chains of inference, involving converging and conflicting arguments. In this section, we show how the ATMS can be used to manage these inferences, keeping track of the tokens on which propositions ultimately depend.

The basic unit of belief in a belief maintenance system is the *belief token*, a special token which carries an attached probability. Belief tokens come in sets. Every extension (maximally specific environment) must contain exactly one token from each set of belief tokens. In the above example, the token V would be paired with another token representing the negation of V. This second token (call it $\overline{V}$) would be assigned belief .4, so that its belief and that of V sum to 1. Belief tokens are processed by the ATMS exactly as are other tokens.

We note here that encoding negations in the ATMS involves defining a *choose* structure and extending the label updating algorithm to include hyperresolution rules. Actually, hyperresolution need not be applied to belief tokens, because the belief computations are performed only on contexts that are maximal with respect to the token sets impacting a proposition's truth value. But deKleer's [1986] disjunction encoding is needed for other exhaustive sets of propositions (e.g., $q$ and its negation $\neg q$).

The belief maintenance system can represent two additional specialized types of token: default and hidden tokens. Default tokens represent propositions that the system chooses to treat as if they were known to be true (i.e., probability 1). Hidden tokens correspond to deKleer's ignored assumptions. They are manipulated by the ATMS, but environments containing them are ignored in the belief computations and are invisible to the problem solver using the belief maintenance system.

Beliefs are computed conditional on the current set of defaults (see below). A belief token may be defaulted (as when a strongly supported hypothesis is provisionally accepted). This causes the token to be treated as if it had probability 1. The other belief tokens in its token set become hidden. The default may subsequently be removed, in which case probabilities on tokens in the set revert to their former values.

Let us return to the inference rule $r \wedge V \rightarrow q$. If $r$ is observed (i.e., declared as a premise), this rule fires and adds the environment {V} to the label of $q$. Absent other rules affecting $q$, the belief in $q$ is equal to the probability of V, or .6. If V is declared as a default, this belief becomes 1.

Now suppose a report from another source indicates the negation of $q$, and that this source is judged to have reliability .8. As before, this can be encoded as an inference rule $s \wedge W \rightarrow \neg q$, where $s$ stands for the source's report and W is a belief token with probability .8. The label of $q$ remains unchanged, but the environment {W} is added to the label of $\neg q$. The tokens V and W imply inconsistent propositions and cannot both be true. The ATMS represents this inconsistency by a *nogood* environment {V,W}.

Belief in $q$ and in $\neg q$ remain unchanged at 1 and 0, respectively. This is so because the label of $q$, the environment {V}, is a default and so has probability 1. The label of $\neg q$ contains the environment {W}. Belief in this label is conditioned on the current defaults, and because W is inconsistent with V its conditional probability given V is zero.

What happened to our belief of .8 in W? Given the default token V, we would have assigned the environment {V,W} probability .8 *if it weren't inconsistent*. But because beliefs are conditioned on the consistent environments, this environment, despite its high prior probability, is discarded as impossible. The probabilities of the consistent environments are divided by .2 so they will sum to 1 after discarding the inconsistent environment.

We see that the belief assigned to inconsistent environments under the current defaults can be thought of as measuring *conflict* associated with the defaults. In this case, we might decide that .8 is too high a degree of conflict, and respond by dropping the default V. Returning to our former belief assignment, the prior probability of the *nogood* environment {V,W} is reduced to .48. Below we list the propositions of interest, the contexts in which they can be proven to hold, and their beliefs.

| Proposition | Context | Belief |
|---|---|---|
| $q$ | $\{V,\bar{W}\}$ | .12/.52 = .23 |
| $\neg q$ | $\{\bar{V},W\}$ | .32/.52 = .62 |

These beliefs are the same as obtained by using each inference rule to define a belief function and combining them by Dempster's Rule.

Rules may have as antecedents the consequents of other rules. For example, consider a third rule $q \wedge X \rightarrow t$. Firing this rule adds the environment {V,X} to the label of *t*, indicating that *t* is true in any context containing both V and X. The probability of this environment, conditional on the defaults (if any) and the consistent environments, defines the degree of belief in *t*. The ATMS automatically keeps track of nonindependencies. For example, there might be another path of reasoning from *q* to *t*. When these rules fire, a second environment containing V will be added to the label of *t*. The probability calculus described below automatically avoids "double counting" the impact of V.

### 3.2 Computing Beliefs

The probability handler computes beliefs on nodes from the probabilities assigned to belief tokens, given the current set of default tokens. The following conditions are assumed:

1. *Each belief token is part of a mutually exclusive and exhaustive set of belief tokens whose probabilities sum to 1.* A set of such tokens corresponds to Shafer's background frame, and carries the basic probability for a belief function.

2. *Each belief token is probabilistically independent of all other belief tokens except other tokens in the same exhaustive set.* This condition may seem restrictive, but in fact it is not. It merely requires that all nonindependencies be represented *explicitly* as shared information (i.e., the labels of two dependent propositions contain belief tokens in the same hypothesis set).

The *belief* in a node is defined as the *probability* of its label. But this probability must be conditioned on the current defaults, and on the consistency of the label. For example, if *x* had label {A,B} where Pr(A)=.8 and Pr(B)=.7, then (absent other information) *x* has belief .56, the product of these probabilities. But suppose *nogood*{B,C}. This means that the conjunction of B and C is impossible, so no belief can be assigned to it. Belief in *x* must thus be revised to reflect this. In other words, belief in *x* is the *conditional* probability of $A \wedge B$, given $\neg(B \wedge C)$:

$$\mathrm{Bel}(x) = \Pr[A \wedge B \mid \neg(B \wedge C)] = \frac{\Pr(\{A,B,\bar{C}\})}{1 - \Pr(\{B,C\})} .$$

If Pr(C) = .8, this expression evaluates to .25. If C is a default, this evaluates to 0 (because the numerator is the probability of an environment that cannot hold under the default).

In general, the belief in a proposition is given by

$$\mathrm{Bel}(node) = \Pr(label \mid defaults, \neg nogood)$$

$$= \frac{\Pr(label \cap \neg nogood \mid defaults)}{\Pr(\neg nogood \mid defaults)} . \quad (*)$$

A simple algorithm for calculating the belief in a proposition (i.e., the probability that it can be proven) follows.

1. Select all environments in the label that are *consistent* with current defaults and that contain no hidden tokens.

2. Remove default tokens from the environments containing them (this amounts to treating them as if they had probability 1). The set of tokens thus defined will be referred to as the *selected tokens*.

3. Select all *nogood* environments that are consistent with current defaults, that contain a selected token, and that contain no hidden tokens. Remove all default tokens and add the remaining tokens to the selected tokens.

4. Repeat Step 3 until no more tokens are added to the selected tokens.

5. The *selected token sets* are those to which the selected tokens belong (e.g., if V is a selected token, the corresponding token set is $\{V,\bar{V}\}$). Form a list of *maximally specific* environments from the selected token sets. That is, form all possible combinations of tokens, taking one from each selected token set.

6. Remove all *nogood* maximally specific environments (i.e., those containing a *nogood* environment, or an environment that is *nogood* when coupled with one or more default tokens). Of the remaining

environments in the maximally specific list, the ones containing a label environment (or which do when combined with one or more default tokens) are those contributing to belief in the node. Add up the probabilities of these (where the probability of the environment is the product of the probabilities of its constituent assumptions) to get the numerator of (*). Add up the probabilities of the whole list (except the removed *nogoods*) to get the denominator of (*). (An alternate way to compute the denominator is 1 minus the probability of the removed *nogoods*.)

This algorithm may be modified to simplify processing of labels in which environments have little overlap [Laskey and Lehner, 1987]. Other efficiency modifications are possible, such as caching intermediate products so that belief computation after label changes can be done incrementally. Although the basic algorithm above works regardless of the pattern of inferential links, the improvement gained by these efficiency modifications will be greatest when nonindependencies are few.

D'Ambrosio [to appear] has suggested an approach very similar to the one described here. His algorithm differs from ours in that his encoding of belief functions is less general, his treatment of *nogoods* that indirectly impact a proposition differs from the network Shafer-Dempster model, and there is no provision for defaults. In the most general networks, both algorithms are exponential in the number of hypothesis sets in the system. D'Ambrosio puts forward some suggestions (not yet implemented) for decreasing complexity; we are currently exploring others.

## 4.0 DISCUSSION

Belief maintenance represents a way of implementing a unified approach to uncertainty management. Any Shafer-Dempster or probabilistic inference network can be represented using this formalism. Indeed, Shafer-Dempster belief theory and belief maintenance without defaults are formally equivalent [Laskey and Lehner, 1987]. The addition of defaults extends Shafer-Dempster theory to include symbolic non-monotonic reasoning. Defaults may be used to represent working assumptions about how to apply the calculus (such as assuming the truth of propositions with a high degree of belief). A by-product of the belief computation is the prior degree of belief assigned to inconsistent hypotheses, which measures the degree of conflict associated with the current defaults. A high degree of conflict can indicate the need to examine the defaults for possible revision.

## REFERENCES


[Cohen, Laskey and Ulvila, 1987] Cohen, M.S., Laskey, K.B., and Ulvila, J.W. *The management of uncertainty in intelligence data: A self-reconciling evidential database* (Technical Report 87-8). Falls Church, VA: Decision Science Consortium, Inc., June 1987.

[D'Ambrosio, to appear] D'Ambrosio B. A hybrid approach to uncertainty, *International Journal of Approximate Reasoning*, to appear.

[deKleer, 1986] deKleer, J. An assumption-based truth maintenance system; Extending the ATMS; Problem Solving with the ATMS, *Artificial Intelligence*, 1986, *28*, 127-162; 163-196; 197-223.

[Doyle, 1979] Doyle, J. A truth maintenance system, *Artificial Intelligence 12*(3), 1979, 231-272.

[Laskey and Lehner, 1987] Laskey, K.B. and Lehner, P.E. Assumptions, beliefs and probabilities, submitted to *Artificial Intelligence*, 1987.

[Shafer, 1976] Shafer, G. *A Mathematical Theory of Evidence*, Princeton, NJ: Princeton University Press, 1976.

[Shafer, 1987] Shafer, G. Belief functions and possibility measures, in Bezdek, J.E. (Ed.), *The analysis of Fuzzy Information*, Boca Raton, FL: CRC Press, 1987.