

PANEL ON DEALING WITH UNCERTAINTY

chairman
Randall Davis
MIT
Cambridge, MA, USA

Participants

Shigetoko Kaihara
University of Tokyo Hospital
Tokyo, Japan

Edward Shortliffe
Stanford University Medical Center
Stanford, CA, USA

Lofti Zadeh
University of California
Berkeley, CA, USA

Introduction

The problem of dealing with uncertainty arises in many areas of AI research and has been attacked with a number of mechanisms. This panel is intended to bring together several of the researchers who have had direct experience with the issue, to report briefly on their efforts and discuss the relative merits of the various approaches.

The panel, originally entitled "Reasoning with Uncertainty", was renamed to encourage a broader view of the topic. "Reasoning", i.e., inference, in the face of uncertainty is only one of the manifestations of the problem and may as a result be too narrow a perspective. Part of the emphasis in the discussion will be to examine the utility of a broader view.

The problem

In the few domains where concepts and inferences can be totally formalized (e.g., pure mathematics), there is no uncertainty of the sort we are discussing here. Unfortunately, in all real problems attacked by AI, we eventually run afoul of the issue. The problem commonly has two sources. First, the basic data we have to deal with may be uncertain. This may occur due to noise (as in speech or vision), due to the fallibility of our measurement procedures (as in medical laboratory

tests), or due to the inherently ill-defined nature of the concept involved (as in "block A is 'near' block B"). Second, the processing applied to that data may introduce its own uncertainty, typically because we lack a comprehensive understanding of the domain. In speech understanding, for instance, we have as yet no theory which allows us to proceed unerringly from waveforms to phonemes to words. At best we can offer plausible guesses at each stage, and must hedge our bets by making multiple such guesses.

The Questions

To provide a focus for the discussion, the participants were asked to consider each of the following questions.

Construing the issue of uncertainty in the broadest sense, what kinds of problems attacked in AI appear to have dealing with uncertainty as a truly central issue? How is uncertainty manifested in those problem areas? For example, work on deductive inference [6], modelling some aspects of common sense reasoning [9], hypothesis formation [5], and signal understanding [1] all had to deal with uncertainty of slightly different form, and all developed different mechanisms.

The psychological connection: What kind of behavior do people display when dealing with uncertainty? What can we learn from this and ought we to try to mimic it in our programs?

What are the current interesting and unsolved issues in representing and dealing with uncertainty?

What would constitute a good solution to the problem? That is, what performance characteristics would it display?

Has there been too much attention paid to the numbers? All of the schemes developed so far eventually use some numeric representation of certainty, but in some the computation and propagation of those numbers is the central concern of the scheme. As we note below, it might be argued that this is an undesirable situation.

Of these questions, the last appears to be potentially the most controversial; it might prove interesting to focus on it in more detail.

As inspiration, consider the following argument that strong emphasis on the numeric aspects of the problem is misguided. First (perhaps not surprisingly) none of the schemes developed to date has a precise notion of uncertainty, a situation which is manifested in the ill-defined nature of the numbers involved

and the processing applied to them. At best, we are given rough guidelines for what the numbers are to mean and some vague assurances that they can be assigned consistently. The arithmetic then done on them is typically ad hoc and its formal properties (e.g., its relation to other models of inference) often not well understood. The overall effect -- ad hoc processing of ill-defined numbers -- does not inspire confidence in the result.

One common response is that not much attention is paid to the precise result, only the range it falls in or the overall ordering it produces. This brings up yet another problem, a by now traditional dilemma: if the systems built around these mechanisms are too sensitive to the numbers produced, that's bad because no one is ready to defend the precise numbers. Yet if the systems are relatively insensitive to the numbers, then all of that mechanism appears spurious. What is the "appropriate" level of sensitivity?

In more general terms, I suggest as a point for discussion the following claim: The numbers and numeric calculations involved in representing and dealing with uncertainty are at best artifacts. Any time the main focus of attention is on the numbers, that's a symptom of something wrong, most likely that there is too little power in the underlying approach to the problem. As a contrast, consider the HEARSAY-II [1] work on speech understanding, in which dealing with uncertainty was viewed as an issue of problem solving, and in which the system architecture and problem solving paradigm were designed to confront the issue. Various measures of uncertainty were introduced (e.g., the credibility of hypotheses, the focus of attention evaluations, etc.), but they played a far less central role because the basic approach -- which included the use of multiple cooperating knowledge sources, best-first search, etc., -- was itself capable of reasonable performance in the face of uncertainty.

As a specific example of the difference in methodology, consider HEARSAY's use of word hypothesis credibility ratings. One way to attempt to get high performance would have been to put effort into developing a finely tuned rating scheme which could be relied upon to give the best score to the correct hypothesis. But this was not done. Instead a version of best-first search was developed which first sought out multi-word hypotheses and then used these as "islands of reliability" from which to press forward the search for a good global hypothesis [2]. The power of the system came from its basic problem solving paradigm, not an elaborate arithmetic model of uncertainty.

Other mechanisms with similar abilities to deal with uncertainty have been explored in other problems:

partial matching, in which the result of the match is not just success or failure, but a symbolic indication of how closely the match was achieved and what parts failed to match;

"constraint sharing", resolving uncertainty by taking advantage of ways in which partial solutions to the problem mutually constrain each other. That is, rather than attempting to resolve every ambiguity as it arises by carefully ranking every alternative, the program continues processing and uses other results to rule out some possibilities. The technique has its

roots in relaxation techniques developed for solving differential equations, but has since been applied in symbolic form to a range of AI problems (e.g., [8,3,7,4]).

The success of such mechanisms on interesting problems and their deemphasis of numeric propagation schemes is, or should be, food for thought.

References

- [1] Erman L, Lesser V, A multi-level organization for problem solving using many diverse cooperating sources of knowledge, *Proc 4th IJCAI*, pp. 483-490.
- [2] Lesser V, et al, Selection of word islands in the HEARSAY-II speech understanding system, *Proc 1977 IEE Conf on ASSP*.
- [3] Lesser V, Corkill D, The application of AI techniques to cooperative, distributed problem solving, Technical Report, COINS Dept, UMass at Amherst.
- [4] Mackworth A, Consistency in networks of relations, TR75-3, CS Dept, Univ. of British Columbia, July 1975.
- [5] Pople H, The formation of composite hypotheses, *Proc 5th IJCAI*, pp 1030-1037.
- [6] Shortliffe E, Buchanan B, A model of inexact reasoning in medicine, *Math Biosci*, 23,351-379(1975).
- [7] Steele G, Sussman G, Constraints, AI Memo 502, Nov 1978, MIT.
- [8] Waltz D, Understanding line drawings of scenes with shadows, in *Psychology of Computer Vision*, Winston, ed.
- [9] Zadeh L, Fuzzy sets, *Info Control*, pp. 338-353 (1965).