The treatment of uncertainty in AI: Is there a better vantage point?

A rejoinder by ALESSANDRO SAFFIOTTI
Dipartimento di Informatica, Università di Pisa, C.so Italia 40, I-56100 Pisa, Italy

The paper I wrote for the *Knowledge Engineering Review* was intended primarily to give a balanced review of different techniques developed in (or imported into) Artificial Intelligence to deal with uncertain knowledge.

History, however, teaches us that every description of facts is inevitably informed by a theory (or opinion) according to which facts themselves are selected, interpreted and presented (Carr, 1961). This was of course also the case with my review, consciously aimed at "highlighting the complex nature of uncertainty and arguing for an open-minded attitude towards its representation and use". This has resulted in what has been called "an eclectic position", the focus of the present debate. The contributions to this debate have in many cases adhered to the "eclectic" position, and given me new hints on the placement of an adequate vantage point from which to survey the question of uncertainty representation.

This contribution to the debate will be structured into two parts: in the first, I will comment on the "eclectic" position, trying to emphasise why I believe this is an "AI view of the treatment of uncertainty", and to clarify the "adequate vantage point" I discussed in my review. In the second part I will give short replies to the commentaries.

How much AI is the "eclectic position"?

My initial interest in the representation of uncertainty in AI sprang from the attempt to merge non-monotonic reasoning and uncertain reasoning. The goal was the definition of a domain independent system which, moving from the Belief Maintenance Systems tradition, was able to operate an "approximate belief revision" on a set of uncertain beliefs. In this light, I approached the uncertainty management arena as a spectator, rather than as a gladiator.

My first impression of the uncertainty arena was an uneasy chaos, and a tendency to attempt resolution by organizing tournaments among the various uncertainty treatment techniques. This feeling remained as I became more acquainted with many of these techniques. The simple and intuitive idea that uncertainty of knowledge is (meta)knowledge and should be treated as such, seemed to have been scarcely taken into account. What I encountered was another instance of the long-lasting AI problem of recognizing, representing and managing a particular kind of knowledge. Accordingly, I saw the various uncertainty treatment techniques as just alternative tools to manage this meta-knowledge. This naturally led me to a position of non-commitment, where the appropriateness of each technique is assessed relative to the task for which an uncertainty management technique is required.

How did this interpretation of uncertainty influence the view presented in my review? The underlying hypothesis in my paper (at least for the "parallel certainty inference" techniques) was that uncertainty be represented together with the knowledge it refers to in agent's knowledge base (KB), and managed by an appropriate technique, like knowledge is represented in the KB and managed by inferencing techniques. Figure 1 will help clarify this hypothesis.[1]

[1] The view of uncertainty discussed here has been made concrete in a general architectural framework for an uncertain agent, detailed elsewhere (Saffiotti, 1988).

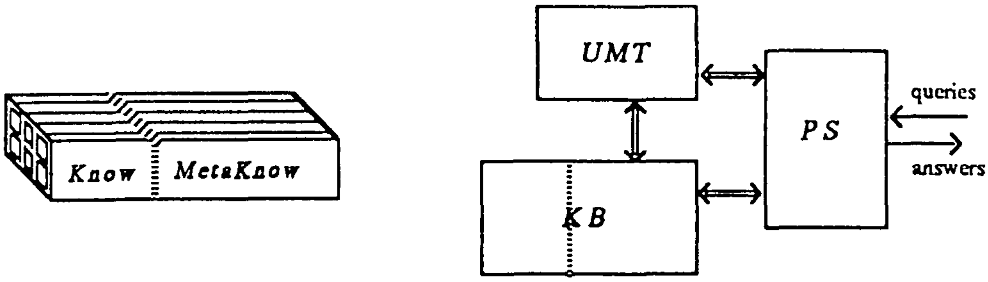

(a) (b)

Figure 1

KB will be in general built out of "bricks" (a), each consisting of an item of knowledge together with some metaknowledge about it. To exemplify, in the vignette of sect. 4.1 of my paper, the rule Mr. Pouring and Mr. Rain are talking about would be represented in the *Know* field, while information about its certainty and its logical dependencies with other items of knowledge would be put in the *MetaKnow* field. As illustrated in (*b*), the uncertainty management technique (UMT) manipulates the *MetaKnow* fields (usually inferring new representations of certainties, exploiting—also—the logical dependencies), while the Problem Solver (PS) manipulates the *Know* fields (usually inferring new representations of knowledge, creating—also—new logical dependencies).

The idea that uncertainty is meta-knowledge has two opposite sides. On the one hand, it is a very weak characterization of uncertainty: it really does not tell us too much about what uncertainty is and how is it to be dealt with. This would not be the case were one to assert that uncertainty is—say—probability, where one would inherit *tout-court* a complete well defined normative theory. On the other hand, this characterization gives us a very powerful view of uncertainty management: accepted thus, uncertainty management can be conceived in a number of different ways, and still be dealt with in the same uniform framework. To see this, it is perhaps interesting to note how this framework can easily accommodate the control approach which Cohen emphasises: the Problem Solver is in fact fully entitled to analyze the meta-knowledge in the KB, and to make use of it when deciding "what to do next". In this extreme case, the UMT module will be probably empty, and the uncertainty present will be inspected and exploited by the rules which control the inference process.[2] What I believe is the important point, from an AI perspective, with this framework, is the ability to exactly place uncertainty in the context of a general intelligent system.

About the contributions

Several interesting points have emerged from the commentaries. Many of the contributors revealed an acceptance of the "eclectic" position, though somewhat different from my own; I hope I have made my point of view a little clearer in the above. In this section, I will deal with some of the points I find worth further discussion; on the points where I am silent, it is understood that I agree. In particular, I am grateful for all the interesting enrichments to my consciously partial review.

John Fox completely supports the "eclectic position", correctly claiming that "numbers, . . . like all other representations . . . are good for some things and less good for others". Nevertheless, in my opinion, his point of view should move even higher: if we agree that uncertainty is knowledge, we must also notice that:

1. It seems to be reductive to circumscribe the debate to a "qualitative vs. quantitative" dichotomy: many other facets and dichotomies can be identified (the well known "intrinsic vagüeness vs. missing information" dicohotomy, for instance).
2. One must keep in mind that different types of uncertainty can in general co-occur in the same problem: so, it is important to distinguish between the task of matching a particular problem to

[2] Or by any other agent's module entitled to inspect the context of the KB.

an uncertainty representation, and that of matching an uncertainty type to an uncertainty management technique.

As for the second point, it is worth noticing that there is a growing interest in the Knowledge Representation community towards the so called "Hybrid Knowledge Representation Systems", in which essentially different types of knowledge can be dealt with by distinct mechanisms in a common frame work (see for example Brachman et al., 1985). Returning to our field, we can quote Zadeh's contribution:

"The three facets described above are not independent and may occur in various guises and different combinations."

This important point leads us to Spiegelhalter's highly stimulating contribution. He seems to suggest a hybrid "number + source" representation as the "most explicit" one. Perhaps not surprisingly, I am going to argue once again that to me the adequateness—and even *explicitness*—of a representation can only be evaluated with respect to a particular uncertainty type (even if the range of adequateness of an hybrid representation is surely much wider than that of a single one). For instance, the use of a number to quantify uncertainty may not be the most sensible choice when the nature of the domain knowledge makes it difficult to extract the numbers we need (whatever their semantics) from experts: the vignette of set.4.1 was intended to emphasise the fact that often the available knowledge is not in the form of—or easily put into—numerical formalisms (though I agree with Spiegelhalter that numerical methods may be adequate when the required numbers are actually available in the object knowledge).

Another interesting issue highlighted by Spiegelhalter is the relationship between "vagueness' of a certainty measure (I prefer not to commit myself to probability values) and their sensitivity to further experience. This actually is an issue I am presently working on, and I don't intend to propose a definite answer; though, I am fairly convinced not only of the opportunity to keep the two concepts distinct, but even that value sensitivity may in general be different with respect to confirmations and to disconfirmations. To illustrate this point, let us consider Church's Thesis, according to which the class of recursively[3] computable functions and that of intuitively computable functions are the same; this Thesis is usually regarded as "highly probable", but no number of confirmations of this conjecture will increase its truth valuation; on the contrary, a single counterexample would be enough to completely invalidate the Thesis.

John Lemmer presents the problem of assessing Dempster-Shafer and Bayes representations in a manner sensitive to knowledge engineering considerations. While trying to claim again that an assessment can only be fully drawn when the application context is taken into account, I should note that I do not fully agree with Lemmer's arguments on the bulk of data required by the two representations. It seems to me that, in many cases, the number of subsets in the frame of discernment to which the probability mass function assigns a non-zero value will be small. This results in a double advantage: by one side, the data we need to extract from the expert will be small as well; by the other side, well known methods can be used to exploit this characteristic and reduce computational burden. On the other hand, Lemmer apparently dodges the problem of the high number of conditional probabilities needed by Bayesian methods.

Philippe Smets describes an interpretation of Dempster-Shafer theory other than the one described in my paper. Unfortunately, he does not dwell on which problems or which uncertainty types his interpretation is best suited to, and to which it is not. I guess that, from an AI viewpoint, the availability of a well defined decision mechanism in Smet's proposal is a most interesting feature, and it is perhaps worth deeper inspection. I also regard as an interesting "trick" the assignment of a non-null mass to the null set, even if I believe that other methods can be employed to deal with the problem of a non exhaustive set of hypotheses (e.g. introducing a dummy "other-possibilities" hypothesis into the frame of discernment).

Some of the contributors opened their own window overlooking the uncertainty forest, and

[3] Or formally computable, according to some other equivalent formalism.

describe the different views they can appreciate from their vantage points; looking at the same phenomenon from different viewpoints is—I believe—an important step towards a better understanding of its shape and dimensions. In addition, Dubois and Prade positively enrich their view by placing the possibility theory into the forest. Though, it is not clear to me how the problem of the choice of an adequate representation looks like from their point of view, and in particular what is the role possibility theory plays in this problem. To this respect, I would like to call Zadeh's observation that "the mere fact that method A is, in principle, a special case of method B does not imply that A is subsumed by B". Another window is opened by Hamid Berenji. From there, we can see the distinction between fuzziness and ambiguity standing in the foreground. The observations about Fox's position (points 1 and 2 above) apply to this distinction as well.

The comments of Paul Cohen are worth deeper inspection. Cohen presents what may be called "the control approach" (briefly described in sect.5.3 of my original paper). In this approach, uncertainty is not necessarily internally represented and manipulated as such, but its presence is taken into account by the control strategy of the system. I have tried to show above how this approach can be accommodated within an unifying framework; to this respect, the control approach can be seen as just one of the many approaches to dealing with uncertainty. As such, it is good for some tasks, less good for others. It is in this light that I believe that "selecting the right calculus" is a most important issue. From this perspective, Cohen's zealous advocacy needs to be challenged: otherwise, Cohen's position could be easily charged with typicality in its smuggling another "best" theory, instead of accepting it as just "one of the available theories".

Another point is worth discussing: Cohen says that "the updating mechanism . . . is pivotal to any theory of updating, but it plays a much smaller role in AI systems". I would agree with him, were I to accept that "the goal of AI is to have computers solve problems". However, I don't feel that sentenctious about the goal of AI. While it may be true that we don't need an internal representation of uncertainty if we have to build artificial agents whose sole goal is "to solve problems" *alone*, I believe that we do need an internal representation of uncertainty if we want our agents to cooperate (even to solve problems *together*) with other agents. Artificial agents, as natural ones, must be able to communicate with one another about their doubts, their uncertainties, as well as about their convictions, and to explain them, if they are to seek exchanges and cooperation. In this case, it will not be true that "a theory of uncertainty . . . should be less concerned with how much one believes propositions, and more concerned with what one does in uncertain situations"; in the case of an autonomous agent who has to solve problems, it probably will.[4] I do not see any good reason to look at one of the these cases as "more AI" than the other.

I have already referred to Zadeh's contribution above. Here I would like to notice how his contribution seems to adhere in a most constructive way to the "eclectic" suggestion, by inspecting some of the facets of the uncertainty phenomenon, and by giving some examples of how to perform the uncertainty facets $\leftrightarrow$ technique matching. This actually corresponds to taking seriously the claim made in my paper that uncertainty theories and uncertainty types are to be seen as "two rows of dancers", and that "our aim would then be to match these dancers, and so find the designated partner technique for the uncertainty type affecting our problem", far from any attempt to set any ordering. I am sure, and I am glad for this, that Zadeh and myself do not "see in different light" this important issue.

Finally, in a personal letter—not reprinted above—, Dr. Dempster kindly pointed out that my statement concerning the paternity of Dempster-Shafer theory (first sentence of sect.3.1) is wrong. The key original reference for the theory is actually (Dempster, 1967), which mainly addresses the finite domain case. To Dempster, I apologize.

Conclusions

I would like to conclude this short contribution by reporting a sentence borrowed from the debate about "which representation for which problem" ubiquitous in the knowledge representation area.

[4] But notice that even the human user of the system may be one of the cooperating agents.

The history of human science and culture illustrates the point that very often progress in some fields depends on the creation of a specific new formalism, with the right epistemological and heuristic power. The same has to be said about formalisms for use in artificially intelligent systems. We need criteria for evaluating formalisms in the light of the uses to which they are to be put.

(Sloman, 1985)

Additional references

Brachman, R, Pigman Gilbert, V and Levesque, H, 1985. "An essential hybrid reasoning system: knowledge and symbol level accounts of KRYPTON", in *Proc. of 9th IJCAI* (Los Angeles, CA).

Carr, E H, 1961. *What is history?* (Macmillan & Co. Ltd., London).

Dempster, A, 1967. "Upper and Lower Probabilities Induced by a Multivalued Mapping", *Annals of Mathematical Statistics* **38**(2).

Saffiotti, A, 1988. "Belief Systems for Doubtful Agents", Technical Report DL-NLP-88-7, University of Pisa (Pisa, Italy).

Sloman, A, 1985. "Why We Need MAny Knowledge Representation Formalisms", in: M. Bramer (Ed.) *Research and Development in Expert Systems*, Cambridge University Press (Cambridge, MA).