Uniform Accountability for Multiple Modes of Reasoning*

Laveen Kanal† and Donald Perlis‡
Computer Science Department, University of Maryland

ABSTRACT

This article discusses various issues surrounding the general debate on knowledge representation methods and argues in favor of the idea that while many methods are necessary, there also must be a kind of unified nature to the enterprise if it is to serve the needs of intelligence. Specific points include the misleading distinction between probabilistic and logical reasoning regarding the notion of truth and also some matters of nonmonotonicity. An effort is made to sample the broad range of approaches in the literature, with an eye toward such a unification.

KEYWORDS: ***commonsense reasoning, artificial intelligence, uncertainty representation, probability theory, probabilistic logic***

INTRODUCTION

The need for a highly expressive language in artificial intelligence, particularly in formalizations of commonsense reasoning, is widely recognized. That such a language will necessarily incorporate many styles of reasoning seems unarguable. In particular, reasoning about uncertainty is very important. Many people have been discussing these issues for decades, and the literature has become quite sophisticated. Particularly troublesome is the relationship of formal treatments to real-world prediction, which also is surely a large part of any test for appropriate commonsense reasoning. The situation is complex,

* This paper is based on a much shorter commentary (Kanal and Perlis, [1]) that critiques Cheeseman [2]. We will occasionally refer to his paper, but our points now are much broader than simply a rejoinder to his stance.

† Supported in part by NSF grants ECS-83-00799 and DCR-8504011 and in part by the L.N.K. Corporation.

‡ Supported in part by grants from the Army Research Office (DAAG29-85-K-0177) and the Martin Marietta Corporation.

Address correspondence to Laveen Kanal, Computer Science Department, University of Maryland, College Park, Maryland 20742.

International Journal of Approximate Reasoning 1988; 2:233–246

52 Vanderbilt Ave., New York, NY 10017 0888-613X/88/$3.50

involving ramifications of logic, probabilities, logical probability, probabilistic logic, subjective probability, and some 46,656 different varieties of Bayesians (Good [3]).

It appears to be agreed that many approaches are important in dealing with complex information. However, there is a need to integrate results from these approaches. Thus, for example, temporal information and probabilistic information both may bear on a problem, and special-purpose temporal and probabilistic reasoning engines may be used. However, each will provide "answers" within its special framework. This leaves us with the meta-problem of interpreting such answers in a unified context relevant to the original problem.

Suppose, for instance, that the probability of being involved in a car accident in 1980 can be calculated with respect to certain priors. This does not mean it is the same as the current probability (in 1988) given today's priors. A temporal reasoner may be used to keep track of "now" in an appropriate way, while a probability engine calculates the "odds" given whatever priors. If an answer of 7.5% is given to some probability query, it must be tagged with additional information as to what the temporal context is, and finally a meta-reasoner (a uniform accounter) must decide what significance to attribute to the various results. Thus it may or may not be deemed appropriate to use a 1980 probability in 1988. This again can lead to a probabilistic calculation as to the chance that the 1980 figure is a good fit to 1988, but note that it can be performed only on the basis of accounting for the temporal as well as probabilistic aspects. In particular, it would be strained and of dubious benefit to assign a probability to the statement that 1980 occurred before 1988.

One way to view this example is that knowledge about priors is itself a sophisticated matter, independent of calculational methods for probabilities. Just how one decides what the priors for a given situation are, such as whether or not they represent current priors or 1980 priors (and how 1980 priors relate to 1988 priors), are issues not in themselves settled by simply calculating a probability. The whole picture must be assessed, in ways not currently understood but clearly in ways that will have to make use of various kinds of information. That is, uniform accountability is needed in order for there to be cooperation between methods, and this in turn requires a language for assessing them.

This language need not be used by each method, however. Viewing the different methods as the petals of a flower, we postulate a central pistil that serves both as communication bus and as arbiter between the petals. The central pistil must represent all methods in its own language. Of course, it is in principle possible to envision a hierarchy of ever-more-encompassing languages, accounting for wider and wider ranges of reasoning modes, with no one superlanguage encompassing all. However, in that case, there would be a theoretical union to this hierarchy that could serve as a kind of universal language; moreover, sufficiently "high" languages in the hierarchy would, from the perspective of any small number of more modest modes, seem highly general. Thus, in any event, a very ambitious level of general accountability appears essential to any kind of reasoning organized toward unified goals in an intelligent way. Otherwise we simply have various petals (modes), none of which knows what the others are doing.

PROBABILITY, LOGIC, AND UNCERTAIN REASONING IN AI

Earlier in the history of AI, methods of approximate or uncertain reasoning including probabilistic and fuzzy-set approaches were largely ignored. The situation is no longer so bleak. For example, recent literature includes Kyburg [4], Kanal and Lemmer [5], Nutter [6], and Halpern and Rabin [7]. Recently, some have advanced probability as the one framework best suited for handling reasoning in the context of uncertainty (Cheeseman [2]). But the issues regarding probability as a particular means for addressing uncertainty are not cut-and-dried. The foundational status of probability theory is very much in debate. Some good references to the many sides of this are given by Fine [8], Good [3], Renyi [9], and a most interesting but apparently not well known book, *Knowing and Guessing*, by Watanabe [10].

In commonsense reasoning we are often not in a position to insist that a statement is absolutely true, but rather that it carries a degree of uncertainty. Now this is widely accepted and the basis for much research in AI from many different directions. There are on-going attempts to model this both with and without the explicit use of numbers to measure uncertainty. However, it is instructive to pursue the suggestion to use probability. A statement such as that "the probability of Fx given Bx in context C is p" is a statement of (presumed) fact (about F, B, C, and p) and is easily written as

$$\text{prob}(`Fx', `Bx \,\&\, C') = p$$

(with quasi quote marks to ensure that propositional or sentential constructs Fx and Bx are properly recorded as terms; see Perlis [11, 12] for details of quoting in regard to beliefs). A more general form is as follows:

$$(*) \qquad p \leq \text{prob}(`Fx', `Bx \,\&\, C') \leq q$$

But the probability statement (*) itself is then apparently one we are being urged to take seriously, as true (about F, B, C, p, and q). Thus (*) is a kind of axiom that is then subject to ordinary (logical) modes of inference. For instance, presumably all would agree that it would be quite awkward to have both (*) and, say,

$$(**) \qquad q + 0.3 \leq \text{prob}(`Fx', `Bx \,\&\, C') \leq 0.9$$

in the same reasoning system (where, say, $q < 0.6$), simply because they are logically contradictory. There is then an underlying arena of full (logical) truth, even in probabilistic reasoning.

Now one might acknowledge this while denying that the internal expressions such as Fx have definite truth values in common sense. But there are several objections that can be made to this. First, if the form of the internal expressions such as Fx allows them to take the form of a probability statement itself, such as (*), then what justification is there for saying (*) is true? Second, the conditions Bx and C are used as if true. Third, Bayesian alternatives seem to presuppose a fixed set of primitive (atomic) notions underlying all of common sense. Fourth, truth does not get in the way of attempts to model nonmonotonicity in logical terms; rather these attempts aim at avoiding the qualification problem of impossibly large numbers of special-case assertions, and a probability approach will suffer the same problem. Fifth, there are clear examples of commonsense statements that are intended to be taken as absolutely true. Sixth, reasoning about *actions* seems to demand a combination of probabilities and outright truth assertions. We will examine each of these in turn, especially with regard to our concern for uniform accountability. First, however, in the next two sections we attempt to outline approaches to modeling inquiry in general, and present some models for structure and uncertainty. Then we amplify on the six points above. Finally, we summarize our views.

APPROACHES TO MODELING INQUIRY

Churchman [13] has categorized approaches to the design of inquiring systems in terms of some underlying philosophical bases. Thus he speaks of Leibnizian, Lockean, Kantian, Hegelian, and Singerian modes of inquiry. He and his former students (Mitroff and Turoff [14]) have discussed the relationships and appropriateness of the different modes of inquiry to well-structured (i.e., information-rich and theoretically understood) and ill-structured problem domains and also to limited versus open-ended objectives of inquiry.

A key virtue of Churchman's categorization is that it helps us see that different modes of inquiry and different models are more natural or useful in different problem domains, depending on our knowledge and on our view of where truth is likely to reside. Thus if we think that "truth is in the theory" and a good model is at hand, the Leibnizian model of inquiry, much used in the physical sciences, is most appropriate. In current parlance, this might be termed a top-down or model-directed approach, wherein the hypotheses are defined a priori and the data being sought are predicted from the model. On the other hand, if in our view, "truth is in the data," then the Lockean mode, which characterizes statistics and exploratory data analysis, is perhaps more appropriate, leading to bottom-up or data-driven methodologies.

In practice, one more often encounters the Kantian mode of inquiry, wherein "truth is partially in the theory and partially in the data," and one pursues a model-directed, data-confirmed, data-driven, model-confirmed feedback approach to refining and selecting among closely competing models that may characterize a problem.

The Leibnizian, Lockean, and Kantian modes of inquiry represent much of what is done in the physical and social sciences and in engineering and statistics. The Hegelian or dialectical mode of inquiry is more appropriate to ill-structured or less-controlled problem domains. Here one's world view determines how to interpret a body of data in order to get information. Thus "truth" emerges from a logical clash between a thesis and an antithesis, leading to a synthesis. Examples from problem domains such as management, economics, law, and politics are easy to cite. The point is that there are many approaches to inquiry, and, depending on the problem domain, one may be more useful than another. A similar statement can be made about modeling structure and about modeling uncertainty. Should one then try to use one approach, namely, probability, to model all uncertainty?

STOCHASTIC AND NONSTOCHASTIC MODELS FOR STRUCTURE AND UNCERTAINTY

Numerous models have been proposed for generating and describing stochastic, nonstochastic, and mixed structures and for representing uncertainty within a given context. It would take us too far afield to attempt to do justice to the vast literature on such work. Here we simply cite some of the many topics that have been addressed, with a few references where additional information on such models and their relationships may be found.

In addition to the references to Fine [8], Good [3], Watanbe [10], and others previously cited, numerous references to the literature on logic, probabilities, logical probability, probabilistic logic, subjective probability, nonnumeric measures of likelihood, fuzzy logic, and the foundations of probability may be found in three volumes of *Uncertainty in Artificial Intelligence* (Kanal and Lemmer [5, 15] Levitt et al. [16]).

Of increasing recent interest are models for stochastic structure known as the Markov mesh and Markov random field models (Kanal [17], Geman and Geman [18]) and network models for probabilistic inference termed causal Bayesian nets (Pearl [19]). Computer science and artificial intelligence have contributed some very novel grammar and graph models for representing structuring, including problem-reduction models, Petri nets, attributed grammars, and semantic nets (see Shapiro [20]). Then, of course, there are various forms of logic (e.g., first-order, modal, temporal—again, see Shapiro [20]). The potential combinations of these various approaches are clearly numerous.

It should perhaps be apparent that just as no single model of inquiry is well suited to or natural for all problem contexts, no single characterization of uncertainty is likely to be suitable for all the many contexts of generative and descriptive models of stochastic and nonstochastic structure. Thus it is not surprising that several approaches to modeling uncertainty have been proposed and many surveys have been written on this subject (for example, Bhatnagar and Kanal [21], Berenstein et al. [22]). Nevertheless there appears to be a school of thought that the only way to model uncertainty is through probability. We are reminded of the aphorism: If the only tool you possess is a hammer, then the whole world looks like a nail.

This is not to say that a probabilistic approach to uncertain reasoning is not valid in many contexts. Spieglhalter [23], in a section titled "Probability—Is it appropriate, necessary, and practical?" and in subsequent sections, makes a very good case for using probabilistic reasoning, in particular a subjectivist Bayesian approach, in characterizing predictive expert systems. Certainly, attempts to relate various ad hoc uncertainty calculi to established formal methods of reasoning such as probability need always to be encouraged. But to argue that probability is the best way to model uncertainty in all contexts is to ignore many aspects of reasoning, some of which are brought out in the following sections.

EPISTEMIC ISSUES AND ACCOUNTABILITY

Regarding the first of the six "objectives" we raised earlier, we ask whether the entire expression (*) is to be regarded as a belief of reasoner g, or is Fx a p-to-q-degree belief of g if Bx & C represents all g knows? It appears that at least some expressions of the form (*) must be available to g for reasoning. And these presumably are taken at face value. For instance, suppose an explanation of a previous choice is needed: Why did g choose A over B? Perhaps because it was more likely to achieve g's goals. In drawing such a conclusion, g reasons *about* the probabilities! So the information that a reasoner needs, even a probabilistic reasoner, includes the probability numbers and even statements about those numbers. Of course, g might assign a degree to these as well, but this leads us into an infinite regress: Either we stop somewhere with an expression g is willing to use at face value, or we perpetually force g back to self-doubt of any expression whatsoever. Thus, in order for g to be able to account for its own reasoning, its belief set must contain meta-information about how it reached conclusions. Thus the variety of methods employed must allow a uniform representation, even if in detail they differ markedly. As another example, if $P(\text{flying } x \mid \text{bird } x) = 90\%$, doesn't this strongly suggest that $(\exists x)(\text{bird } x \ \& \ \neg \text{flying } x)$ (as true, and not merely possible)? And conversely, is it not the existence of nonflying birds that makes the probability statement useful? This seems to say that both probabilistic and nonprobabilistic statements are necessary and that they are closely interrelated.

Thus probabilistic information should be made to fit neatly into a reasoning system so that logical consequences can be drawn from the probabilities. What is needed is a graceful accountability for g's reasoning, so that the (undeniably many and complex) modes of reasoning can be related to one another by g, and further conclusions by g can be based on g's analysis of what g has thought. This is hard; we contend that it is one of the major stumbling blocks in efforts to capture commonsense reasoning. Yet much more effort seems to be expended in perfecting individual modes and arguing their primacy than in looking for flexible blends. The need for accountability suggests that whatever procedures and whatever assertional representations are available, *some* aspect of the reasoning must have uniform access to a large portion of these so that it can make general judgments about the reasoning behavior as a whole. Otherwise there would seem to be no possibility of our agent g responding to the question "Why did you do that?" with something like "I did that because I thought X-Y-Z; and it was a mistake; next time I will do U-V-W." That is, the various methods employed by $g(X, Y, Z)$ must be at least somewhat accessible to looking backward and to interrelatability sufficient to allow intelligent conclusions. The language in which this occurs need not be that in which each separate reasoning method occurs; but the accountability language must be able to encode adequate portions or summaries of the others.

In particular, uncertainties and certainties should satisfy such a requirement. Thus one area for research is the graceful incorporation of numerical probabilities and nonnumerical approaches to both uncertain and certain (traditional) modes of reasoning. We will explore some issues in the interrelation of (traditional) logic and (numerical) probability in what follows, as a preliminary step in that direction. We note that logic might be an interesting candidate language of accountability; this in no way would mean that other methods should be constrained to a formal logical form, but simply that suitable summary features of those methods should be encodable in it. As one further indication of what we have in mind, consider that algorithms used in computer vision do not resemble theorem-proving or formal logic; and yet the *results* of such algorithms must be representable in declarative form for us to reason about them—for example, to decide to take action from having seen a missile approaching. Note that this particular suggestion strongly resembles procedural attachment (Nilsson [24]) and is contested by Pentland and Fischler [25].

TRUTH AND BELIEF

Now to the second objection, that the very use of Bayesian probabilities relies on a sense of truth of basic propositions. For consider Cheeseman's [2] conclusion that "I might decide to throw out the milk based on the probability value (95%) using the information [that it has been in the refrigerator for three days and there is a bad smell]." Now, clearly he has decided that *this* is the information to use (to regard as true about the real-world context) rather than his other scenario (in which there is no bad smell and the probability is then only 1%). So whenever probabilities actually are used, one assesses what conditional assertions actually are to be trusted for the given situation, which amounts to regarding those conditions as representing the reasoner's unconditional beliefs as to the truth. For if the reasoner hedges and attributes further probabilities here, then the computation gets pushed back until at some point a final stance is taken.

Now, one may argue that there is no problem with true primitive (atomic) assertions but rather that it is quantified logical conditionals that are not to be accorded full truth in commonsense settings. Thus perhaps it is only well-formed formulas (wff's) such as $(\forall x)(Px \rightarrow Qx)$ that are never really true or false but only contingent. But here there are clear counter examples. We will present these below in our discussion of point five.

PRIMITIVES OF COMMON SENSE

This brings us to our third objection—namely, that there seems to be no fixed set of primitive or atomic concepts that underlie all of commonsense reasoning. To be sure, some have argued otherwise (e.g., Schank [26]). But in order to employ Bayes' theorem, we first need to be given the prior probabilities, and the calculation from these of some further probability had better not conflict with already given information. Thus a successful use of the Bayesian approach depends on assuming as valid certain prior probabilities or distributions or empirical data, and then compound probabilities may be calculated from these. But if instead we are given various compound probabilities, there is no guarantee of consistency with Bayes. Moreover, even if there were an effective way to guarantee this, it would not be very useful for commonsense reasoning unless Schank turns out to be right about a set of universal primitives, and this appears unlikely. Many commonsense concepts are intertwined rather than defined in a neat hierarchy.

NONMONOTONICITY

Our fourth objection has to do with nonmonotonicity. Finding the whole of what is known about F, as is required in determining the context C, amounts to precisely the usual default problem. How is the totality of relevant things determined? The brute-force approach is to enumerate all special cases, such as that each of a very large number of conceivable situations is not to be taken as relevant. This is impractical and amounts roughly to the frame problem or the associated qualification problem. The logical approach has at least addressed this issue, and circumscription even has a computationally attractive handle on it.

Thus relevance is, in a sense, *the* issue in nonmonotonicity, and it is unclear that probability has much to offer on this matter.

For instance, if P(flies | bird) = 90%, and we seek P(flies | bird & ostrich) and we don't know anything else, how can we decide that ostrich is irrelevant so that still P = 90%? That is, if we know nothing about ostriches (not even that they are birds), it is important to recognize that nothing is known and that therefore a default should be relied upon. Bayesian theory will not help us out; rather it will leave us high and dry with a question—What are the values of prob(flies | ostrich), prob(bird | flies & ostrich), prob(bird | ostrich)?—even if we know ostriches *are* birds, so that the latter two numbers are 100%.[1] We still get no useful answer. One might want to assume that an ostrich is a very typical bird, in the absence of contrary information; but what mechanism does this? In the absence of contrary information, we (by default) make certain probability assumptions.

One might argue that g should avoid any conclusion if certain probabilities are missing (e.g., for ostriches). But this is surely no good. If instead of Ostrich, we were told Blinked-eye(x), that x blinked its eye, should we also refuse to calculate the chance that the bird flies? Maybe only ostriches have eyelids, in which case, in fact, x will not fly. But we do not know this. As far as we know, it is a totally irrelevant bit of data. How do we decide the irrelevance? It is not merely a matter of *searching* the whole database as Pearl says; it is worse. We must determine that no conclusion as to the flying status of x is provable from anything in the database; that is, we must decide that flying is independent of eye blinking and the rest of hte data. Now this is in general an undecidable problem, and in special cases it is being addressed head-on by the various nonmonotonic formalisms. Cohen [28] and Grosof [29] consider this point.

DEFINITIONAL BELIEFS

We now comment on our fifth objection. Consider definitional beliefs, such as democracy is a form of government, or cabbage is a kind of vegetable. These are not mere typicality (or plausible but uncertain) statements, and they (and countless more like them) are surely part of everyday commonsense reasoning. For a more detailed example, if Bob is a state employee, and if state employees

[1] This is nontrivial, however. As Pearl has noted [27], the material implication $P \rightarrow Q$ does not really correspond to the (single) conditional probability prob($Qx | Px$) = 1, due to the implicit presence of *all* conditioning information in the latter; that is, it is assumed that Px is *all* one knows when Qx is concluded from the given statement. This is the source of the supposed nonmonotonic component available in probabilistic reasoning. However, one can simulate the material implication as follows by the use of *many* conditional probabilities: $(\forall X)$(prob($Q | P \& X$) = 1). This says that *no matter what* other information is present, the presence of P necessitates that of Q. Thus, there is a precise connection that can be expressed formally, but it is a complex one.

are not eligible to win the state lottery, then Bob is (literally and absolutely) not eligible to win the state lottery. This is, as far as we can see, a simple matter of ordinary reasoning that does not involve probabilities and that would be strained if couched in terms of probabilities. Again, if a student asks "are all carcinogens related to cancer?" the answer is (by definition) "Yes!"

TAKING ACTION

This still leaves us with the issue of when to actually go with a particular belief, in terms of taking action. Here presumably we want not only that things be weighed in various terms, but also that directives be given—for example, if X then DO Y. Here DO can be a predicate symbol whose truth value is connected to, say, physical events that occur when DO is proved. The form of X may even be probabilistic, such as

$$\text{prob}(Z, \text{all-I-know}) = .9$$

Moreover, often the best plan is to seek more information before taking further action. For this it again is necessary to recognize one's ignorance. Suppose you want to catch a plane, and you know that on the average planes depart about 30 minutes late, say with probability 95% of being at least 15 minutes late. Are you going to plan to arrive late at the airport in the expectation that the particular plane you want will be late, or will you call the airport to verify this? What if the probability had been 99.99%? What about 50%? There seems to be a clear need for judgments as to information-gathering action, based on some kind of cutoffs. To be sure, there is a well-known formalism for this kind of situation, namely, decision theory. However, decision theory makes use of assumed losses, values, or utilities, which themselves are not statistical or probabilistic, and the calculated outcomes will generally lead to nonrandomized decisions. So decision theory is already something partly outside of probability theory.

A probability is not an end in itself; it is information that can serve to guide further thought and action. One might, on the basis of information that a glass has fallen and that falling glasses often break, decide to ask (or find out) whether indeed it has broken, or at least to consider that it might have broken. Despite frequently expressed belief to the contrary, classical logic does not fly in the face of this. Indeed, given the axiom Fallen(glass), neither Broken(glass) nor $\neg$Broken(glass) follows, and in fact the given axiom is perfectly consistent with the additional axiom Possible('Broken(glass)'). This is very much in the spirit of McCarthy's characterization of circumscription as a rule of conjecture, and was the point of Perlis [30], where in fact the latter wff is *derivable* on the basis of the given axiom (alone). The truth orientation of logic is not tied to any particular view of the world, and certainly not to one of guaranteeing certainties about matters of empirical fact. The user of a logic has to design it via suitable choice of predicates and axioms, including statements of likelihood and uncertainty if desired. See Nutter [6] and Halpern and Rabin [7] for recent work on this.

SUMMARY AND CONCLUSIONS

Pearl [31] gives a convincing use of probability in solving the Hanks-McDermott shooting problem [32]. However, there are by now many solutions to this problem; among the most impressive, to our thinking, are those of Lifschitz [33] and Haugh [34], which use ordinary circumscription. Thus once again various methods show themselves to be useful, and arguing that any one is necessarily the right one seems counterproductive. The mark of intelligence should be the ability to look at things in many ways and assimilate them into an informed and rounded viewpoint. Humans easily go back and forth between multiple analyses of problems, recognizing that here is a probabilistic insight, there a default, and over there a deductive consequence. When we encounter someone who has trouble shifting mental gears to incorporate another view, our response is to regard them as demonstrating a certain lack of intelligence in that area.

We have claimed that truths are part of commonsense reasoning. This does not mean that we think *all* wff's are necessarily to be taken as true or as false, of course. We grant that many commonsense assertions have associated degrees of acceptance and truth; but this is already widely acknowledged in AI. We are sympathetic in particular to the contention that probabilistic thinking has a role in commonsense reasoning. For example, Spiegelhalter's case for using subjective probabilities in expert systems that are subject to numerical evaluation is quite convincing. There are other (engineering) examples such as optical character recognition, where a Bayesian probabilistic approach is possible to an important part of the problem but is very awkward in comparison to a simple nonprobabilistic structural approach (Kanal and Chandrasekaran [35]). It is necessary to demonstrate the natural contexts and advantages for different approaches with clear detailed examples based on problems taken from the literature to afford comparison among approaches, and even more important to allow effective work on combining approaches flexibly. Approaches for a general accountability language have been suggested at various times. One is that of predicate calculus, strongly urged by Nilsson [24] in terms of procedural attachment. It is true that logic is a very expressive medium, and we have qualified sympathy with this choice. It is also plausible, however, that for certain purposes other means of expression may be more useful or that new forms of "logic" will have to be invented to adequately accommodate a broad range of reasoning modes.

As noted earlier, AI, and more generally computer science, have shown a lot of creativity in developing models for nonstochastic structures. In like manner there has been considerable creativity in stochastic modeling. The modeling of commonsense reasoning will take even more creativity, requiring more than the current approaches to probabilities and logic. We suggest that it is important to come up with plausible and detailed solutions to problems requiring multimodal approaches, with graceful accountability allowing reasoned conclusions from assessments of past reasoning. Thus the variety of reasoning methods must fit into a kind of unification, a currently elusive effort.

References

1. Kanal, L., and Perlis, D., Discussion of P. Cheeseman's paper, *Comput. Intelligence,* 1988, to appear.
2. Cheeseman, P., Inquiry into computer understanding, *Comput. Intelligence*, 1988.
3. Good, I. J., *Good Thinking: The Foundations of Probability and Its Applications*, Univ. of Minnesota Press, Minneapolis, 1983.
4. Kyburg, H., Bayesian and non-Bayesian evidential updating, *AI* **31**, 271–293, 1987.
5. Kanal, L. N., and Lemmer, J. F. (Eds.), *Uncertainty in Artificial Intelligence*, North-Holland, Amsterdam, 1986.
6. Nutter, J. T., Default reasoning using monotonic logic: a modest proposal, *Proceedings AAAI*, 297–300, 1983.
7. Halpern, J., and Rabin, M., A logic to reason about likelihood, *AI* **32**, 379–405, 1987.
8. Fine, T., *Theories of Probability*, Academic Press, New York, 1973.
9. Renyi, A., *Foundations of Probability*, Holden-Day, San Francisco, 1970.
10. Watanabe, S. [1969] *Knowing and guessing*, Wiley, New York.
11. Perlis, D., Languages with self-reference I, *AI* **25**, 301–322, 1985.
12. Perlis, D., Languages with self reference II, *AI*, to appear.
13. Churchman, C. W., *The Design of Inquiring Systems*, Basic Books, New York, 1971.
14. Mitroff, I. I., and Turoff, M., The why's behind the how's, *IEEE Spectrum*, March, 62–71, 1973.
15. Lemmer, J. F., and Kanal, L. N. (Eds.), *Uncertainty in Artificial Intelligence*, Vol. 2, North-Holland, Amsterdam, 1988.
16. Levitt, T., Kanal, L., and Lemmer, J. (Eds.), *Uncertainty in Artificial Intelligence*, Vol. 3, North-Holland, Amsterdam, to appear.

17. Kanal, L., Markov mesh models, in *Image Modeling* (A. Rosenfeld, Ed., Academic Press, New York, 239–243, 1980.

18. Geman, S., and Geman, D., Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images, *IEEE Trans. Pattern Anal. Mach. Intell.* **6**, 721–741, 1984.

19. Pearl, J., Fusion, propagation, and structuring in Bayesian networks, *AI* **29**, 241–288, 1986.

20. Shapiro, S. C. (Ed.), *Encyclopedia of Artificial Intelligence* (2 vols.), Wiley-Interscience, New York, 1987.

21. Bhatnagar, R., and Kanal, L., Handling uncertain information: a review of numeric and non-numeric methods. In *Uncertainty in Artificial Intelligence* (L. N. Kanal and J. F. Lemmer, (Eds.), North-Holland, Amsterdam, 3–26, 1986.

22. Berenstein, C., Kanal, L., and Lavine, D., Consensus and evidence. In *Pattern Recognition in Practice*, Vol. 2 (E. S. Gelsema and L. N. Kanal, (Eds.), North-Holland, Amsterdam, 1986.

23. Spiegelhalter, D. J., A statistical view of uncertainty in expert systems, in *Artificial Intelligence*, (W. Gale, (Ed.), Addison-Wesley, Boston, 1986.

24. Nilsson, N., Artificial intelligence prepares for 2001, *AI Mag.* **4**, 7–14, 1983.

25. Pentland, A., and Fischler, M., A more rational view of logic, or, up against the wall, logic imperialists! *AI Mag.* **4**, 15–18, 1983.

26. Schank, R., *Conceptual Information Processing*, North-Holland, Amsterdam, 1975.

27. Pearl, J., Discussion of P. Cheeseman's paper, *Comput. Intelli.* 1988, to appear.

28. Cohen, M., An expert system framework for non-monotonic reasoning about probabilistic assumptions, in *Uncertainty in Artificial Intelligence*, Vol. 1 (L. N. Kanal and J. F. Lemmer, Eds.), North-Holland, Amsterdam, 279–293, 1986.

29. Grosof, B., Non-monotonicity in probabilistic reasoning, in *Uncertainty in Artificial Intelligence*, Vol. 2 (J. F. Lemmer and L. N. Kanal, Eds.), North-Holland, Amsterdam, 1988, to appear.

30. Perlis, D., Circumscription as introspection, in *Methodologies for Intelligent Systems*, (Z. Ras and M. Zemankova (Eds.), North-Holland, Amsterdam, 440–444, 1987.

31. Pearl, J. A., Probabilistic treatment of the Yale shooting problem, Tech. Report CSD-8700XX R-100, UCLA, September 1987.

32. Hanks, S., and McDermott, D., Default reasoning, nonmonotonic logics, and the frame problem, *AAAI-86*, 328–333, 1986.

33. Lifschitz, V., A formal theory of action, *IJCAI-87*, 966–972, 1987.

34. Haugh, B., Simple causal minimizations for temporal persistence and projection. *AAAI-87*, 218–223, 1987.

35. Kanal, L., and Chandrasekaran, B., On linguistic, statistical and mixed models for pattern recognition, in *Frontiers of Pattern Recognition* (S. Watanabe, Ed.), Academic Press, New York, 163–192, 1969.