Let Many Flowers Bloom: A Response to "An Inquiry into Computer Understanding"*

Joseph Y. Halpern

IBM Almaden Research Center
650 Harry Road
San Jose, CA 95120

email: halpern@almaden.ibm.com

There are two main issues I want to address in this response to Peter Cheeseman's "An inquiry into computer understanding" (hereafter, abbreviated as *Inquiry*) and "In defense of 'An inquiry into computer understanding'" (hereafter, abbreviated as *Defense*): (1) the "right" approach to reasoning about uncertainty and (2) the relationship between logic and probability.[1] Cheeseman's positions on these issues is at times confusing, and occasionally contradictory (as he admits in *Defense*). Rather than belaboring the contradictions, I will try to focus on the high-level issues, which I believe are important ones for AI.

*This paper was written while I was on sabbatical at the University of Toronto. I'd like to thank the Computer Science Department at the University of Toronto for their kind hospitality. This version is essentially identical to one that appears in *Computational Intelligence* **6**, 1990, pp. 184–188.

[1]Both of Cheeseman's articles, as well as a number of responses to them, including the responses of Hayes, Israel, Ruspini, and Schubert mentioned below, appear in *Computational Intelligence*, volume 5, number 1.

1 Approaches to uncertainty

Cheeseman makes a number of strong statements in *Inquiry* and *Defense*. Perhaps the most important of them can be summarized by the following two claims: (1) all the information we have about a situation about which we want to reason can be best represented in terms of probabilities, and (2) if we get new information, we should use Bayes' rule of conditioning to update our earlier probabilities.

Probability is a good method for representing uncertainty. I consider it to be the best (and certainly the best understood!) of all the methods currently available. I suspect it will end up being the method of choice for many applications. Indeed, I believe that many of the other methods proposed for reasoning under uncertainty can be best understood in terms of probability theory. (See, for example, [HF89] for some discussion and further references on the issue of how the Dempster-Shafer approach can be understood in terms of probability theory; see [Gro86, Hec86] for a discussion of how MYCIN's certainty factors can be understood in terms of probability theory.) However, it does not follow that probability should be the method of choice for *every* application (or for all parts of any one application).

Even if one is firmly committed to using probability theory, one does not have to be a Bayesian. There are many problems with the Bayesian position espoused by Cheeseman, mainly stemming from the question of how to assign probabilities and how they can be computed. Since these issues have been discussed at length in the literature, I will just briefly review them here. Suppose we have an estimate of the probability of some event A which we want to update in the light of some new evidence E. In order to use Bayes' rule, we must know the quantities $Pr(A \cap E)$ and $Pr(E)$. However, the event E might have been completely unanticipated, so we might not have assigned a probability to it. Even if we are in a constrained situation, where all the possible items of evidence can be predicted beforehand, there are typically too many possibilities for us to feasibly assign them all probabilities beforehand. Indeed, it is often the case that we have no idea of the prior probability of observing certain items of evidence, so our assignment of prior probabilities (i.e., the probabilities $Pr(E')$ for each possible item of evidence E' that we might observe) can be quite arbitrary.

It may well be that the Bayesian approach can ultimately handle these objections, but I believe that many, if not most, statisticians would dispute Cheeseman's claim that all the problems have been solved. Indeed, in his zeal to convert everyone to the Bayesian approach, Cheeseman may have done a disservice to those researchers trying to make a careful case that it does have wide applicability.

Even if the Bayesian approach can handle all these objections, it does not immediately follow that it is always the *right approach* to use. There is some reason to believe for a number of applications, other approaches can be quite successful. MYCIN, which uses *certainty factors*, has demonstrated performance equal to that of experts [YFW+79]. The Japanese claim that fuzzy logic is being used with great success to control subway trains in the city of Sendai [Pol89]. These are not isolated examples. To me, this suggests an underlying robustness in many of the applications in which we are interested. A closer look at the MYCIN system provides some further evidence. Heckerman shows [Hec86] that there are several restrictions that a system must satisfy in order to guarantee that certainty factors will have a number of reasonable properties. However, these restrictions are often violated in practice. So why does MYCIN work so well? Heckerman suggests that this is because detailed considerations of uncertainty are not critical to the system's performance, citing a sensitivity analysis by Cooper and Clancy [BS84] which reveals that the system's performance did not change significantly when many of the certainty factors in the system were changed.

It seems to me that we need to understand more about the nature of the uncertainty that faces us in real world problems. Is it indeed relatively insensitive to changes in representation? If so, why? Can we somehow exploit this feature? Humans do relatively little quantitative reasoning (and are not so good at the little they do [TK74]). Yet we are remarkably good at commonsense reasoning, which surely involves a lot of reasoning under uncertainty. What approaches are we using?

Cheeseman suggests that we focus all our energies on the Bayesian approach. "Let a single flower bloom", he says in *Defense*. I find this somewhat premature. At this early stage, I think we need more variety in our garden. I suspect that crossbreeding several strains will give us our best hope of success.

2 Logic and probability

From time to time, Cheeseman identifies *logic* with first- or higher-order logic. This is a rather narrow view of logic. Typically, one designs a logic to reason about *something*. Thus, there are logics for reasoning about time, for reasoning about knowledge, for reasoning about obligations, and so on. One standard approach, in both AI and philosophy, is to identify a logic with a set of axioms and rules of inference (or, more accurately, the set of consequences of a given set of axioms and rules of inference). Rather than taking such a proof-theoretic view of logic, I would like to take a different (although still reasonably standard) viewpoint, which is more model theoretic. I will take a logic to consist of a formal language and semantics for formulas in the language.

From this point of view, if we want to use logic to help us in our reasoning, we must construct an appropriate syntax and semantics. There is a great deal of complexity in the real world. Typically we want to reason about certain phenomena of interest, ignoring many other details which are taken to be irrelevant. In order to reason about these phenomena in a rigorous way, we try to capture them using a formal mathematical model. Doing this well is an art. It must be done in such a way so as to make it easy to relate the objects in the semantics to the phenomena we are trying to reason about. Once we have a good idea of what we want to reason about, we must design a formal language in which to carry out that reasoning. The syntax should be well-suited to its purpose. That is, if we want to make statements about probability, it should be easy to render the statements we want to make as formulas in the logic. If we can't capture a lot of statements we want to make, the logic isn't doing its job.[2] Even if the relevant concepts can be captured by the logic, if the resulting formula is awkward and unnatural, then the logic is still not doing its job right: the language does not have the right constructs. Of course, "ease" and "naturalness" will often depend on the beholder; nevertheless, they are goals to strive for.

What's the point in using logic? To me, the main reason is that a formal

[2]Of course, we might be willing to give up some expressive power in order to simplify the language and perhaps cut down the complexity of checking validity of formulas in the language. But if we give up too much expressive power, then there is no point in considering the logic.

logic helps guide our reasoning. A good semantics can help us see relationships between situations much better than symbol manipulation. A good language can help us express and clarify important, but subtle, distinctions. (Note the relationship between this motivation and the desire to have an elegant syntax and a natural semantics. A poor syntax will make it difficult or impossible to capture subtle distinctions; a poor semantics will typically obscure rather than clarify what is going on.)

From this point of view, propositional logic is intended to capture and clarify the subtleties of reasoning about conjunction, negation, and (especially) implication. First-order logic is intended to capture and clarify issues regarding universal and existential quantification and (because it includes propositional logic), issues regarding conjunction, negation, and implication as well. The syntax and semantics of these logics are well-suited for these tasks. They are not particularly well-suited to reasoning about probability.

In *Defense*, Cheeseman discusses a robot Robbie who reasons about, and has beliefs about, the world. He says that "the essential argument of *Inquiry* is that classical [first- or higher-order] logic is an inadequate tool for Robbie to use to represent the world and to reason about it, because it does not allow degrees of belief (all propositions are either true or false), and it does not explicitly condition its conclusions on the evidence on which they are based." I would agree with the first part of Cheeseman's statement. To the extent that it is important to represent probabilistic information (which I believe it is), classical logic is inadequate. (Although I believe it will be an important component of the logic that should ultimately be used, simply because it is important to represent universal and existential information, and information about implication, as well as probabilistic information.)

On the other hand, to the extent that I understand it, I find the second half of the of Cheeseman's statement somewhat confused. In particular, it is certainly possible to allow statements about degrees of belief in the language (such as the statement "the probability of φ is 1/2" or even "Robbie's subjective probability of φ is 1/2") while still having all formulas be either true or false. The comment that classical logic does not explicitly condition its conclusions on the evidence on which they are based seems to be somewhat of a red herring. First-order logic has no notion of evidence. If the claim is simply that in order to represent Robbie's beliefs, it would be useful to have explicitly in the language a notion of conditioning, then I certainly accept that. But I do not accept that there is a problem with the way first-order logic is used to reach conclusions.

One way to deal with inadequacy of first-order logic as a tool to represent Robbie's world is to extend first-order logic to include probability. However, if we plan to do this, then we must deal with a number of the well-known subtleties that need to be captured when it comes to reasoning about probability. One subtlety that arises is that there are many different uses of the English word "probability". Carnap [Car50] focussed on two notions, which he called probability$_1$ and probability$_2$. Probability$_2$ corresponds to *statistical assertions*, such as "the probability that a (randomly chosen) bird flies is .9". This is quite different from a statement about a degree of belief, such as "the probability that Tweety (a particular bird) flies is .1". The latter statement roughly corresponds to Carnap's notion of probability$_1$. Given that these are two different notions, we would like to be able to capture the difference, both semantically and syntactically. I have suggested one way of doing so in [Hal89], expanding on ideas of Bacchus [Bac88]. The idea is that we have two different types of probability functions. In order to capture statistical assertions, we use a probability on the domain. Thus, to check the probability that a randomly chosen bird flies, we need a probability function on the set of birds. In order to capture degrees of belief, we can imagine that there are many *possible worlds*; in some of these worlds, Tweety flies, in others, Tweety doesn't fly. If we put a probability on the set of possible worlds, we can capture a statement about degrees of belief. The probability that Tweety flies is .9 if the set of possible worlds where $Flies(Tweety)$ is true has probability .9. The language enforces these intuitions by having two different symbols, one corresponding to each type of probability.

We can use these ideas to help us formalize the distinction between the following two statements considered by Cheeseman in *Inquiry*:

(A) The probability that something will float given that it is made of plastic is 0.2.

(B) The probability that all plastic things will float is 0.2.

Statement (A) is clearly intended to be a statistical assertion; that is, it seems to be saying that roughly 20% of plastic objects float. Its truth can be checked empirically (at least in principle), by performing experiments on the world. Statement (B) is an assertion about a degree of belief in a particular assertion, namely, the assertion that all plastic things float. Although the assertion is either true or false (in this case, almost certainly false), an agent may have a certain degree of belief in its truth.

Cheeseman argues in *Inquiry* that (A) and (B) should be captured, respectively, by the formulas

(A′) $\forall x Pr(Float(x)|Plastic(x), c) = 0.2$

(B′) $Pr(\forall x(Float(x)|Plastic(x), c)) = 0.2,$

where the c on the right-hand side of the conditional represents the "general context", that is, the sum total of background information that the agent has acquired thus far. These formulas show that Cheeseman does seem to implicitly accept the need for a language in which to capture the distinctions he would like to make. He does not provide a formal syntax and semantics, however, and so loses the discipline provided by having a good logic for reasoning about probability. Since probabilistic statements are subtle, it is not surprising that he falls prey to a number of common errors.

As Schubert points out in his response, Cheeseman's formula (A′) does not capture statement (A), at least, not under the standard interpretation of the universal quantifier in first-order logic. In first-order logic, for any constant $\mathbf{c}$, the formula $(\forall x \varphi(x)) \Rightarrow \varphi(\mathbf{c})$ is valid; we can always instantiate a universally quantified variable. Conversely, if $\varphi(d)$ holds for each element d of the domain, then $\forall x \varphi(x)$ holds. Thus, under the standard interpretation of $\forall$, we can see that Cheeseman's statement amounts to saying that for every particular object in the domain, we believe that the probability that that object floats given that it is plastic is 0.2. Now suppose we take d to be the particular plastic duck currently floating in our bathtub. We surely don't believe that $Pr(Float(d)|Plastic(d), c) = 0.2$. Notice that by replacing the universal quantifier by a particular domain element, we have passed from statistical reasoning to reasoning about degrees of belief. This by itself should suggest that there is something wrong here.

In *Defense*, Cheeseman tries to justify his usage by saying that a substitution such as the one described above is not permitted by Cox's axioms. I frankly don't understand his point. In any case, he does seem to be saying that his notion of universal quantifier is *not* the same as the standard first-order universal quantifier. I completely agree with Schubert that what Cheeseman means to use here is a different sort of a quantifier, essentially meaning "for a randomly chosen x". We can in fact design a language which includes a quantifier $\forall^r$, so that $\forall^r x\varphi(x)$ holds if $\varphi(x)$ holds for a randomly chosen x. (This, in fact, is done in [Bac88, Hal89, Kei85], although the actual syntax used is slightly different.) This new quantifier has some of the flavor of a universal quantifier, but does not share all of its properties. A logic with a good syntax and semantics can make the distinction between the two quantifiers clear.

Now let us consider the formula (B′). To me, it seems meaningless. The problem is that in first-order logic, a formula such as $\forall x\varphi$ is well formed only if φ is a well-formed formula. Unfortunately, in (B′), the argument of the universal quantifier is the expression $(Float(x)|Plastic(x), c)$, which is not a well-formed formula. Since the formula that says all plastic objects float is $\forall x(Plastic(x) \Rightarrow Float(x))$, the formula that seems to best capture what Cheeseman intends by statement (B) is

C′ $Pr(\forall x(Plastic(x) \Rightarrow Float(x))|c) = 0.2.$

At the risk of being repetitive, let me summarize by saying that, given context c, there are three distinct statements that can be considered here:

(A″) the probability that a randomly chosen plastic object floats is 0.2,

(B″) for each plastic object d in the domain, the probability that d floats is 0.2,

(C″) the probability of the assertion "all plastic objects float" is 0.2.

Statement (A″) is a statistical assertion; its truth can be checked by performing experiments on the world. The latter two statements are statements about degrees of belief. Although each plastic object in the world either floats or it doesn't, it still makes sense to say that the degree of belief that an agent has in whether a particular plastic d floats is 0.2. Although Cheeseman intends (A) to be identical to (A″), the formula (A′) actually corresponds to (B″) under the standard interpretation of the universal quantifier. What is needed to capture (A) and (A″) is a different kind of quantifier.

It seems to me that the subtleties here are perhaps the best argument for getting good logics of probability.

A good logic of probability can help clarify another issue discussed at length by Cheeseman and some of the other respondents: subjective vs. objective truth. There is no controversy over the fact that the truth of certain statements may depend on context and on who is uttering the statement. For example, a statement such as "it is raining here now" may be true in Toronto on Feb. 22, 1990 and false in San Jose on Aug. 19, 1989. It is also seems reasonable that the truth of a statement like "I believe that the probability of rain in Toronto on Feb. 22, 1990 is 1/3" depends on who is making the statement. Philosophers and logicians have long been aware of issues involving *indexical* statements such as these [Kap79, Per79]. However, Cheeseman seems to be making the claim that the truth of *all* statements is subjective. The philosphical issues here have been discussed at length in the literature (and were brought up in a number of other responses, including those of Hayes, Israel, Ruspini, and Schubert), so I will not pursue them further here. What I instead want to consider is the more pragmatic issue of how this claim should affect the way we design systems.

Whether or not there is a notion of objective truth, it seems useful when modeling a robot like Robbie to postulate a world (and an associated notion of truth) external to Robbie. For example, in *Inquiry*, Cheeseman considers a situation where some milk in his refrigerator might be "off". Given only the information that the milk has been in the refrigerator for three days, Cheeseman suggests that we ascribe probability 0.1 to *offmilk*. In the light of additional information (such as the observation of a decaying piece of meat in the refrigerator), we might be inclined to change this probability. In any case, by making such an assignment, Cheeseman himself seems to implicitly accept the usefulness of talking about whether the milk is "really" off or not in the real world.

An issue arises as to whether there is a difference between saying something is true in the real world and assigning it probability one. It is not clear precisely what Cheeseman's position is here. In Section 2 of *Defense*, he says "A good heuristic is to regard beliefs with very high probability as 'true' (i.e., their probability = 1)." Later, in Section 4, he states that "in a Bayesian framework ... no finite amount of evidence gives absolute support to an hypothesis." This suggests that no proposition is ever assigned probability 1. In terms of making decisions, it may often be reasonable to treat beliefs held with high probability as true (this can be viewed as a default rule), as long as this can be retracted.[3] Note that such an approach quickly leads to issues of nonmonotonicity (although not necessarily to nonmonotonic logic).

All these issues—in particular, the distinction between truth and probability—can be clarified by using the notion of possible worlds. We can distinguish one world as designating the actual world (where, say, the milk is off), and other worlds as representing the way the world might be. The fact that the milk is off and Robbie believes that it is highly probable that the milk is not off can be captured by a model where in the actual world the milk is off, but the set of worlds where the milk is not off is given high probability by Robbie. The possible-worlds framework allows us to place probability one on a fact that is false; this could be the case in a model where the actual world has probability 0. It is similarly easy to capture another situation discussed by Cheeseman, where most people believe a certain fact and yet it is false.

Notice that in possible-world semantics, truth is context-dependent. A fact may be true at one world and false at another. However, in any fixed world, it is either true or false. In particular, in the world designated to be the actual world, it is either true or false.[4]

[3] However, as Henry Kyburg has observed, such retractions may not always be possible. If you assign high probability to a proposition that your parachute will open, and act on it, then it may be too late to retract your jump once you discover that the parachute is not working.

[4] It is probably worth pointing out here one serious technical error that Cheeseman makes in arguing against an objective notion of truth. Schubert in his article mentions that Goldbach's conjecture (every even number greater than 2 is the sum of two primes) is an example of a statement that is either true or false. Cheeseman's says (in Section 4 of *Defense*) that "Schubert's example of Goldbach's conjecture being either true or false rather than being either strongly believed or strongly disbelieved is just an assertion of faith that a proof or disproof will eventually be discovered, and this proof is not just another case of 'proof' such as the numerous Fermat's last theorem 'proofs'." Cheeseman is quite simply wrong here. He has confused truth and provability, two important but quite distinct notions. Goldbach's conjecture is either true or false of the natural numbers. That is, either every even number greater than 2 is indeed the sum of two primes, or there is a counterexample, some even number that is not the sum of two primes. While most mathematicians believe Goldbach's conjecture to be true, there is not the same strong belief that a proof can be found. Whether or not a fact can be proved depends on the axiom system being used. Gödel's famous incompleteness theorem tells us that for

3 Conclusions

I hope I've made it clear that I agree with Cheeseman that probability theory has an important role to play. On the other hand, it seems to me that we have too little understanding of the issues involved in reasoning about uncertainty in real-world situations to commit to any one approach at this time. I suspect that in many situations, many approaches will work, due to the fundamental robustness of the world we are living in. In other situations (such as commonsense reasoning), we will probably need to use a number of approaches to adequately represent what is going on. Neither classical logic nor pure probability theory will be enough by themselves (as Cheeseman himself admits in Section 6 of *Inquiry*). A great deal more research needs to be done to understand the strengths and weaknesses of the techniques at our disposal.

Whatever technique we use, I believe that a good logic for reasoning about uncertainty will have an important role to play. Indeed, Cheeseman's own writing provides an object lesson in the need for a formal logic. Cheeseman makes no attempt to be precise or formal in his statements about how he wants to deal with probability. While too much formality can often obscure a point, too little may lead to confusion. By using a formal logic along the lines suggested above, with a good syntax and semantics, his statements (or, more accurately, various interpretations of his statements) can be formalized. It is much easier to discuss how reasonable Cheeseman's assumptions are once we have a formal model; it has the added advantage of forcing us to be clear about what the probability is being taken over.

Acknowledgements: I'd like to thank Fahiem Bacchus, Russ Greiner, Henry Kyburg, and particularly Moshe Vardi for their helpful comments on an earlier draft of this paper.

References

[Bac88] F. Bacchus. *Representing and reasoning with probabilistic knowl-*

any reasonable (recursive) axiom system, there will be many true facts about the natural numbers that cannot be proved. They are true nonetheless, and their truth is independent of our beliefs (and is not affected by the fact that they may not be provable).

edge. PhD thesis, University of Alberta, 1988. Also issued as Waterloo University Technical Report CS-88-31.

[BS84] B. G. Buchanan and E. H. Shortliffe, editors. *Rule-based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, 1984.

[Car50] R. Carnap. *Logical Foundations of Probability*. University of Chicago Press, 1950.

[Gro86] B. N. Grosof. Evidential confirmation as transformed probability. In L. N. Kanal and J. F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, pages 153–166. North-Holland, 1986.

[Hal89] J. Y. Halpern. An analysis of first-order logics of probability. In *Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, pages 1375–1381, 1989. A revised and expanded version appears as IBM Research Report RJ 6882, 1989, and will appear in *Artificial Intelligence*.

[Hec86] D. Heckerman. Probabilistic interpretations for MYCIN's certainty factors. In L. N. Kanal and J. F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, pages 167–196. North-Holland, 1986.

[HF89] J. Y. Halpern and R. Fagin. Two views of belief: Belief as generalized probability and belief as evidence. Research Report RJ 7221, IBM, 1989. A revised version will appear in March, 1990.

[Kap79] D. Kaplan. On the logic of demonstratives. In P. A. French, Jr. T. D. Uehling, and H. K. Wettstein, editors, *Contemporary Perspectives in the Philosophy of Language*, pages 401–412. University of Minnesota Press, Minneapolis, 1979.

[Kei85] H. J. Keisler. Probability quantifiers. In J. Barwise and S. Feferman, editors, *Model-Theoretic Logics*, pages 509–556. Springer-Verlag, 1985.

[Per79] J. Perry. The problem of the essential indexical. *Noûs*, 13:3–21, 1979.

[Pol89] A. Pollack. Fuzzy computer theory: How to mimic the mind? *New York Times*, April 2 1989.

[TK74] A. Tversky and D. Kahneman. Judgement under uncertainty: Uncertainty and biases. *Science*, 185:1124–1131, 1974.

[YFW+79] V. L. Yu, L. Fagan, S. M. Wraith, et al. Antimicrobial selection by a computer: A blinded evaluation by disease experts. *Journal of the American Medical Association*, 242(12):1279–1282, 1979.