

LIKELIHOOD, PROBABILITY, AND KNOWLEDGE

Joseph Y. Halpern

IBM Research Laboratory
San Jose, California 95193

David A. McAllester

MIT Artificial Intelligence Laboratory
Cambridge, Massachusetts 02139

Abstract: The modal logic LL was introduced by Halpern and Rabin [HR] as a means of doing qualitative reasoning about likelihood. Here the relationship between LL and probability theory is examined. It is shown that there is a way of translating probability assertions into LL in a sound manner, so that LL in some sense can capture the probabilistic interpretation of likelihood. However, the translation is subtle; several more obvious attempts are shown to lead to inconsistencies. We also extend LL by adding modal operators for knowledge. The propositional version of the resulting logic LLK is shown to have a complete axiomatization and to be decidable in exponential time, provably the best possible.

1. Introduction

Reasoning in the presence of incomplete knowledge plays an important role in many AI expert systems. One way of representing partially constrained situations is with sentences of first-order logic (cf. [MH, Li, Re]). Any set of first-order sentences specifies a set of possible worlds (first-order models). While such assertions can deal with partial knowledge, they cannot adequately represent knowledge about relative likelihood. This problem was noted by McCarthy and Hayes ([MH]), who made the following comments:

We agree that the formalism will eventually have to allow statements about the probabilities of events, but attaching probabilities to all statements has the following objections:

1. It is not clear how to attach probabilities to statements containing quantifiers in such a way that corresponds to the amount of conviction that people have.
2. The information necessary to assign numerical probabilities is not ordinarily available. Therefore, a formalism that required numerical probabilities would be epistemologically inadequate.

There have been proposals for representing likelihood where a numerical estimate, or certainty factor, is assigned to each bit of information and to each conclusion drawn from that information (see [DBS, Sh, Za1] for some examples). But none of these proposals have been able to adequately satisfy the objections raised by McCarthy and Hayes. It is never quite clear where the numerical estimates are coming from; nor do these proposals seem to capture how people approach such reasoning. While people seem quite prepared to give qualitative estimates of likelihood, they are often notoriously unwilling to give precise numerical estimates to outcomes (cf. [SP]).

In [HR], Halpern and Rabin introduce a logic LL for

reasoning about likelihood. LL uses a modal operator L to help capture the notion of "likely", and is designed to allow qualitative reasoning about likelihood without the requirement of assigning precise numerical probabilities to outcomes. Indeed, numerical estimates and probability do not enter anywhere in the syntax or semantics of LL.

Despite the fact that no use is made of numbers, LL is able to capture many properties of likelihood in an intuitively appealing way. For example, consider the following chain of reasoning: if P_1 holds, then it is reasonably likely that P_2 holds, and if P_2 holds, it is reasonably likely that P_3 holds. Hence, if P_1 holds, it is somewhat likely that P_3 holds. (Clearly, the longer the chain, the less confidence we have in the likelihood of the conclusion.) In LL, this essentially becomes "from $P_1 \Rightarrow LP_2$ and $P_2 \Rightarrow LP_3$, conclude $P_1 \Rightarrow L^2P_3$ ". Note that the powers of L denote dilution of likelihood.

One way of understanding likelihood is via probability theory. To quote [HR], "we can think of likely [the modal operator L] as meaning 'with probability greater than α ' (for some user-defined α)". The exact relationship between LL and probability theory is not studied in [HR]. However, a close examination shows that it is not completely straightforward. Indeed, as we show below, if we simply translate "P holds with probability greater than α " by LP, we quickly run into inconsistencies. Nevertheless, we confirm the sentiment in the quote above by showing that there is a way of translating numerical probability statements into LL in such a way that inferences made in LL are *sound* with respect to this interpretation of likelihood. Roughly speaking, this means that if we have a set of probability assertions about a certain domain, translate them (using the suggested translation) into LL, and then reason in LL, any conclusions we draw will be true when interpreted as probability assertions about the domain. However, our translation is somewhat subtle, as is the proof of its soundness; several more obvious attempts fail. These subtleties also shed some light on nonmonotonic reasoning.

We enrich LL by adding modal operators for knowledge, giving us a logic LLK which allows simultaneous reasoning about both knowledge and likelihood. This extends the logics used in [MSH1, Mo]), where knowledge has been treated in an all or nothing way: either a person knows a fact or he doesn't. However, there are many cases in which knowledge is heuristic or probabilistic. For example, suppose I know that Mary is a woman, but I have never met her and therefore do not know how tall she is. Under such circumstances, I consider it unlikely that she is over six feet tall. However, suppose that I am told that she is on the Stanford women's basketball team. My knowledge about her height has now changed, although I still don't know how tall she is. I now consider it reasonably likely that she is over six feet tall.

LLK gives us a convenient formal language for reasoning about such situations.

LLK can be shown to have a complete axiomatization, which is essentially obtained by combining the complete axiomatization of LL with that of the modal logic of knowledge. In addition, we can show that there is a procedure for deciding validity of LLK formulas which runs in deterministic exponential time, the same as that for LL. This is provably the best possible.

In the next section, we review the syntax and semantics of LL. In Section 3, we discuss the translation of English sentences into LL and show that there is a translation which is sound with respect to the probabilistic interpretation of L. In Section 4, we add knowledge to the system to get the logic LLK. Detailed proofs of theorems and further discussion of the points raised here can be found in [HMc], an expanded version of this paper.

2. Syntax and semantics

We briefly review the syntax and semantics of LL. We follow [HR] with one minor modification: for ease of exposition, we omit the "conceivable" relation in the semantics (and thus identify the operator L^* of [HR] with the dual of G). We leave it to the reader to check that all our results will also hold if we reinstate the conceivable relation. The reader should consult [HR] for motivation and more details.

Syntax: Starting with a set $\Phi_0 = \{P, Q, R, \dots\}$ of *primitive propositions*, we build more complicated LL formulas using the propositional connectives \neg and \wedge and the modal operators G and L . Thus, if p and q are formulas, then so are $\neg p$, $(p \wedge q)$, Gp ("necessarily p "), and Lp . We omit parentheses if they are clear from context. We also use the abbreviations $p \vee q$ for $\neg(\neg p \wedge \neg q)$, $p \Rightarrow q$ for $\neg p \vee q$, $p \equiv q$ for $(p \Rightarrow q) \wedge (q \Rightarrow p)$, Fp ("possibly p ") for $\neg G\neg p$, and $L^i p$ for $L \dots L p$ (i L's).

Semantics: We give semantics to LL formulas by means of Kripke structures. An *LL model* is a triple $M = (S, \mathcal{L}, \pi)$, where S is a set of *states*, \mathcal{L} is a reflexive binary relation on S (i.e., for all $s \in S$, we have $(s, s) \in \mathcal{L}$) and $\pi: \Phi_0 \times S \rightarrow \{\text{true}, \text{false}\}$. (Intuitively, π assigns a truth value to each proposition at all the states.)

We can think of (S, \mathcal{L}) as a graph with vertices S and edges \mathcal{L} . If $(s, t) \in \mathcal{L}$ then we say that t is an \mathcal{L} -*successor* of s . Informally, a state s consists of a set of hypotheses that we take to be "true for now". An \mathcal{L} -*successor* of s describes a set of hypotheses that is reasonably likely given our current hypotheses. We will say t is *reachable* (in k steps) from s if, for some finite sequence s_0, \dots, s_k , we have $s_0 = s$, $s_k = t$, and $(s_i, s_{i+1}) \in \mathcal{L}$ for $i < k$.

We define $M, s \models p$, read p is satisfied in state s of model M , by induction on the structure of p :

- $M, s \models P$ for $P \in \Phi_0$ iff $\pi(P, s) = \text{true}$,
- $M, s \models \neg p$ iff $\text{not}(M, s \models p)$,
- $M, s \models p \wedge q$ iff $M, s \models p$ and $M, s \models q$,
- $M, s \models Gp$ iff $M, t \models p$ for all t reachable from s ,
- $M, s \models Lp$ iff $M, t \models p$ for some t with $(s, t) \in \mathcal{L}$.

Definitions: A formula p is *satisfiable* iff for some $M = (S, \mathcal{L}, \pi)$ and some $s \in S$ we have $M, s \models p$; p is *valid* iff for all $M = (S, \mathcal{L}, \pi)$ and all $s \in S$ we have $M, s \models p$. It is easy to check that p is valid iff $\neg p$ is not satisfiable. If Σ is a set of LL formulas, we write $M, s \models \Sigma$ iff $M, s \models p$ for every formula $p \in \Sigma$. Σ *semantically implies* a formula p , written $\Sigma \models p$, if, for every model M and state s in M , we have $M, s \models \Sigma$ implies

$M, s \models p$.

3. The probabilistic interpretation of likelihood


Lp is supposed to represent the notion that " p is reasonably likely". Certainly one way of interpreting this statement is " p holds with probability greater than or equal to α ". However, as already noted in [HR], there are problems with this interpretation of Lp . Suppose we take $\alpha = 1/2$, and consider a situation where we toss a fair coin twice. If P represents "the coin will land heads both times", and Q represents "the coin will land tails both times", then we clearly have $L(P \vee Q)$, as well as $\neg LP$ and $\neg LQ$. But, for any LL model, $L(P \vee Q)$ is true iff $LP \vee LQ$ is true, giving us a contradiction.

We solve this problem by changing the way we translate statements of the form " p is reasonably likely" into LL. Note that if a state s satisfies the formula p (i.e. $M, s \models p$), this does *not* imply that p is necessarily true at s , but simply that p is one of the hypotheses that we are taking to be true at this state. We must use Gp to capture the fact that p is *necessarily* true at s , since $M, s \models Gp$ iff $M, t \models p$ for all t reachable from s , and thus in *no* state reachable from s is $\neg p$ taken to be an hypothesis. The English statement "The coin is likely to land heads twice in a row" is really "It is likely to be (necessarily) the case that the coin lands heads twice in a row", (and not "It is a likely hypothesis that the coin lands heads twice in a row") and thus should be translated into LGP rather than LP . Similarly, "the coin is likely to land tails twice in a row" is LGQ , while "it is likely that the coin lands either heads or tails" is $LG(P \vee Q)$. With these translations, we do not run into the problem described above, for $LG(P \vee Q)$ is not equivalent to $LGP \vee LGQ$. These observations suggest that the only LL formulas which describe real world situations are (Boolean combinations of) formulas of the form $L^i G C$, where C is a Boolean combination of primitive propositions. We will return to this point later.

Having successfully dealt with that problem, we next turn our attention to translating statements of conditional probability: "if P , then it is reasonably likely that Q " or " Q is reasonably likely given P ". The obvious translation of "if P then likely Q " would be $P \Rightarrow LQ$. The arguments of the previous paragraph suggest that we should instead use $GP \Rightarrow LGQ$, but even this translation runs into some problems.

Consider a doctor making a medical diagnosis. His view of the world can be described by primitive propositions which stand for diseases, symptoms, and test results. The relationship between these formulas can be represented by a joint probability distribution, or a Venn diagram where the area of each region indicates its probability, and the basic regions correspond to the primitive propositions.

For example, the following Venn diagram might represent part of the doctor's view, where P_1 and P_2 represent diseases, and P_3 and P_4 represent symptoms:



The diagram shows (among other things) that

- (a) disease P_1 is reasonably likely given symptom P_3 ,
- (b) P_3 is always a symptom of P_2 ,
- (c) if a patient has P_2 , then it is not reasonably likely that he also has P_1 ,
- (d) P_1 and P_4 never occur simultaneously.

The second statement is clearly $G(P_2 \Rightarrow P_3)$ from which we can deduce $GP_2 \Rightarrow GP_3$. Now suppose that we represented the first and third statements, as suggested above, by $GP_3 \Rightarrow LGP_1$ and $GP_2 \Rightarrow \neg LGP_1$, respectively. Then simply using propositional reasoning, we could deduce that $GP_2 \Rightarrow LGP_1 \wedge \neg LGP_1$, surely a contradiction.

The problem is that when we make such English statements as " P_1 is reasonably likely given P_3 " or "the conditional probability of P_3 given P_1 is greater than one half", we are implicitly saying "given P_1 and all else being equal" or "given P_1 and no other information", P_3 is likely. We cannot quite say "given P_1 and no other information" in LL. Indeed, it is not quite clear precisely what this statement means (cf. [HM]). However, we can say "in the absence of any information about the formulas P_1, \dots, P_k which would cause us to conclude otherwise", and this suffices for our applications. In our present example, P_1 is reasonably likely given P_3 as long as we are not given $\neg P_1$ or P_2 or P_4 . Thus, a better translation of " P_1 is reasonably likely given P_3 " is:

$$\neg G \neg P_1 \wedge \neg GP_2 \wedge \neg GP_4 \wedge GP_3 \Rightarrow LGP_1.$$

Similarly, "if a patient has P_2 , then it is unlikely that he has P_1 " can be expressed by:

$$\neg GP_1 \wedge GP_2 \Rightarrow \neg LGP_1.$$

In general, we must put all the necessary caveats into the precondition to avoid contradictions.

This translation seems to avoid the problem mentioned above, but how can we be sure that there are no further problems lurking in the bushes? We now show that, in a precise sense, there are not.

Fix a finite set of primitive propositions $V = \{P_1, \dots, P_n\}$. An atom of V is any conjunction $Q_1 \wedge \dots \wedge Q_n$, where each Q_i is either P_i or $\neg P_i$. Note that there are 2^n such atoms. Let $AT(V)$ be the set of atoms of V , and let $LIT(V) = \{P, \neg P \mid P \in V\}$ be the set of literals of V . We say a function $Pr: AT(V) \rightarrow [0, 1]$ is a probability assignment on V if $\sum_{C \in AT(V)} Pr(C) = 1$. A propositional probability space is a pair $W = (V, Pr)$, where V is a finite set of primitive propositions and Pr is a probability assignment on V .

Let $BC(V)$ consist of all the Boolean combinations of the propositions in V . If $C, D \in BC(V)$, we write $C \leq D$ if $C \Rightarrow D$ is a propositional validity. We extend Pr to $BC(V)$ via $Pr(D) = \sum_{\{C \in AT(V) \mid C \leq D\}} Pr(C)$. If $Pr(D) \neq 0$, we define the conditional probability of C given D , $Pr(C \mid D) = Pr(C \wedge D) / Pr(D)$.

We now consider a restricted class of probability statements about the domain W . Fix α with $0 < \alpha < 1$. A probability assertion about W is a formula in the least set of formulas closed under disjunction and conjunction, and containing conditional probability statements of the form $Pr(C \mid D) \geq \alpha^i$ and $Pr(C \mid D) < \alpha^i$, where $i \geq 0$, $C, D \in BC(V)$ and $Pr(D) > 0$. (Closure under negation is built into these formulas since, for example, $\neg Pr(C \mid D) \geq \alpha^i$ iff $Pr(C \mid D) < \alpha^i$.) Note that by taking $D = \text{true}$, we get $Pr(C) \geq \alpha^i$ or $Pr(C) < \alpha^i$, and by taking $i = 0$ in the former term, we can assert that a certain statement holds with probability one). Since we are dealing with a discrete probability space, this amounts to saying that the statement is true.

Corresponding to these probability assertions about W , we will consider the standard LL formulas over V . These are formed by taking formulas of the form $L^i GC$ and $\neg L^i GC$, $i \geq 0$, where $C \in BC(V)$, and closing off under conjunction and disjunction. By the observations above, these are, in some sense, exactly those LL formulas that describe a "real world" situation involving the primitive propositions of V .

We want to translate probability assertions about W into standard LL formulas over V . As discussed above, a conditional probability assertion of the form $Pr(C \mid D) \geq \alpha^i$ will be translated into a formula of the form $\neg GQ_1 \wedge \dots \wedge \neg GQ_k \wedge GD \Rightarrow L^i GC$, where Q_1, \dots, Q_k are the "necessary caveats". We now make the notion of a "necessary caveat" precise. Given $C, D \in BC(V)$, and $Q \in LIT(V)$, we say Q has negative (resp. positive) impact on C given D in W if

$$\begin{aligned} &Pr(D \wedge Q) > 0 \text{ and } Pr(C \mid D \wedge Q) < Pr(C \mid D) \\ &(\text{resp. } Pr(D \wedge Q) > 0 \text{ and } Pr(C \mid D \wedge Q) > Pr(C \mid D)). \end{aligned}$$

Thus Q has negative (resp. positive) impact on C given D in W if discovering Q lowers (resp. increases) the probability of C given D . We say Q has potential negative (resp. positive) impact on C given D in W if for some $D' \leq D$, Q has negative (resp. positive) impact on C given D' in W . Note that if Q does not have potential negative impact on C given D in W , then once we know D , no matter what extra information we get, finding out Q will not lower the probability that C is true. Similar remarks hold for potential positive impact. We define

$$\begin{aligned} PNI(C, D) &= \{Q \in LIT(V) \mid \\ &Q \text{ has potential negative impact on } C \text{ given } D\}, \\ PPI(C, D) &= \{Q \in LIT(V) \mid \\ &Q \text{ has potential positive impact on } C \text{ given } D\}. \end{aligned}$$

Now using the idea of potential positive and negative impact, we give a translation $q \mapsto q^t$ from probability assertions about W to standard formulas over V . We first define

$$\begin{aligned} [Pr(C \mid D) \geq \alpha^i]^t &= (\bigwedge_{Q \in PNI(C, D)} \neg GQ) \wedge GD \Rightarrow L^i GC, \\ [Pr(C \mid D) < \alpha^i]^t &= (\bigwedge_{Q \in PPI(C, D)} GQ) \wedge GD \Rightarrow \neg L^i GC, \end{aligned}$$

and then translate conjunctions and disjunctions in the obvious way; i.e., if p, q are probability assertions about W , then $(p \vee q)^t = p^t \vee q^t$ and $(p \wedge q)^t = p^t \wedge q^t$. Again we note that the term $\bigwedge_{Q \in PNI(C, D)} Q$ (resp. $\bigwedge_{Q \in PPI(C, D)} Q$) in the translation of $Pr(C \mid D) \geq \alpha^i$ (resp. $Pr(C \mid D) < \alpha^i$) is intended to capture the idea of "putting in all the necessary caveats in order to avoid contradictions".

We now consider a family of translations Tr_D , $D \in BC(V)$, from standard LL formulas over V to probability assertions about W . Roughly speaking, we want $L^i GC$ to be translated to $Pr(C \mid D) \geq \alpha^i$. This will be the effect of Tr_{true} . Using Tr_D relativizes everything to D ; we require this greater generality for technical reasons. Let

$$\begin{aligned} Tr_D(L^i GC) &= Pr(C \mid D) \geq \alpha^i, i \geq 0, \\ Tr_D(\neg L^i GC) &= Pr(C \mid D) < \alpha^i, i \geq 0. \end{aligned}$$

Again, conjunctions and disjunctions are translated in the obvious way, so that if p, q are standard LL formulas:

$$\begin{aligned} Tr_D(p \vee q) &= Tr_D(p) \vee Tr_D(q) \text{ and} \\ Tr_D(p \wedge q) &= Tr_D(p) \wedge Tr_D(q). \end{aligned}$$

Finally, let $CON(W) = \{C \in BC(V) \mid C \text{ is a conjunction of formulas in } LIT(V) \text{ and } Pr(C) > 0\}$. (We take the empty conjunction to be true; of course, $Pr(\text{true}) = 1$.)

With these definitions in hand, we can now state the theorem which asserts that there is a translation from probability assertions about W into LL which is sound.

Theorem 1: Let Σ be a set of probability assertions true of W , and Σ^t the result of translating these formulas into LL (via $p \rightarrow p^t$). If q is a standard LL formula which is semantically implied by Σ^t (i.e., $\Sigma^t \models q$), then for all $D \in \text{CON}(W)$, $\text{Tr}_D(q)$ is a probability assertion true about W .

The theorem follows from two lemmas, which are proved in [HM_c]. The first shows the relationship between the translations described above.

Lemma 1: If q is a probability assertion true of W , then $\text{Tr}_D(q^t)$ is true of W for all $D \in \text{CON}(W)$.

(We remark that neither Lemma 1 nor Theorem 1 holds for arbitrary $D \in \text{BC}(V)$ (a counterexample is given in [HM_c]). Since we are mainly interested in Tr_{true} , this point will not greatly concern us here, but it is interesting to note that we could have modified the translation $p \rightarrow p^t$ so that Theorem 1 did hold for all $D \in \text{BC}(V)$ with $\text{Pr}(D) > 0$. The idea would be to allow PNI and PPI to include arbitrary elements of $\text{BC}(V)$, rather than just literals. The cost of doing this is that the translation could be doubly exponential in the size of V , rather than just linear. If for some reason we are interested in Tr_D for $D \in \text{BC}(V)$, another (less expensive) solution to the problem is to add a new primitive proposition Q to V , extend Pr so that $\text{Pr}(Q=D)=1$, and consider Tr_Q instead.)

We next construct an LL model $M_W = (S, \mathcal{E}, \pi)$ corresponding to the propositional probability space W . The set of states S consists of countably many copies of each $C \in \text{BC}(V)$ with $\text{Pr}(C) > 0$. Successive copies are connected by \mathcal{E} , as well as a state you are likely to move to as your knowledge increases. More formally,

$$S = \{(C_i | i \geq 0, C \in \text{BC}(V), \text{Pr}(C) > 0)\},$$

$$\mathcal{E} = \{(C_i, C_i), (C_i, C_{i+1}) | i \geq 0\} \cup$$

$$\{(C_i, D_0) | D \leq C, \text{Pr}(D | C) \geq \alpha^{i+1}\}.$$

The definition of π is somewhat arbitrary. All we require is that $M, C_i \models C$, for all $C_i \in S$. For definiteness, we define π as follows. For each $C \in \text{BC}(V)$ such that $\text{Pr}(C) > 0$, choose some atom $D \in \text{CON}(W) \cap \text{AT}(V)$ such that $D \leq C$ (such a D must exist since $\text{Pr}(C) > 0$). Call this atom $\text{AT}(C)$. Then $\pi(P, C_i) = \text{true}$ iff $\text{AT}(C) \leq P$. We leave it to the reader to check that with this definition, $M, C_i \models C$.

The following lemma relates truth in M_W to truth in W .

Lemma 2: If q is a standard LL formula, then $M_W, C_0 \models q$ iff $\text{Tr}_C(q)$ is true of W .

Proof of Theorem 1: Suppose Σ is a set of probability assertions true of W , M_W is the canonical model for W constructed above, q is a standard LL formula over V such that $\Sigma^t \models q$, and $D \in \text{CON}(W)$. By Lemma 1, for each formula $p \in \Sigma$, we know that $\text{Tr}_D(p^t)$ is true of W . By Lemma 2, it now follows that $M_W, D_0 \models p^t$. Thus $M_W, D_0 \models \Sigma^t$. Since $\Sigma^t \models q$, we also have $M_W, D_0 \models q$. By another application of Lemma 2, it follows that $\text{Tr}_D(q)$ is true of W . \square

Discussion of the theorem: Theorem 1 shows that by putting in all the "necessary caveats", we do indeed get a sound translation. But in a real world situation, it is not always possible to compute PNI(C, C') or PPI(C, C'), either because we may not know whether a given literal Q should be in one of these sets, or because the set of primitive propositions V may be so large that the computation is impractical. Indeed, in the examples discussed in [McC], V is viewed as being essentially infinite. If we take P to be "Tweety is a bird" and Q to be "Tweety can fly", then Q is likely given P as long as Tweety is not an ostrich, Tweety is not a penguin, Tweety is not dead, Tweety's wings are not clipped, The list of possible disclaimers is endless.

Our assumption of having only finitely many primitive propositions does seem to be both epistemologically and practically reasonable in many natural applications. For example, in medical diagnosis we could take V to consist of relevant symptoms, diseases, and possible treatments, where the symptoms are qualitative (his temperature is very high) rather than quantitative (his temperature is 104° F.).

In any case, if we cannot compute PNI or PPI, and instead use a subset in the translation, then our reasoning may be unsound (in the sense of Theorem 1). This may help to explain where the nonmonotonicity comes from in certain natural language situations. People often use a type of informal default reasoning, saying "P is likely given Q", without specifying the situations where the default Q may not obtain. Of course, this means that the conclusion Q may occasionally have to be withdrawn in light of further evidence. If, on the other hand, we "play it safe", by replacing PNI(C, C') (resp. PPI(C, C')) wherever it occurs in the translation by a superset, it is straightforward to modify to proof of Theorem 1 to show that the resulting translation is still sound.

We have viewed Theorem 1 as a soundness result. It is natural to ask if there is also a complementary completeness result. For example, suppose q is a standard LL formula over V , and for all propositional probability spaces $W = (\text{Pr}, V)$, and all $D \in \text{CON}(W)$, we have $\text{Tr}_D(q)$ true of W for all choices of α in the translation. Is it then the case that q is a valid LL formula? Unfortunately, the answer is no.

To see this, first note that

$$\text{Tr}_D(\text{LGP} \vee \text{LG}\neg\text{P}) = \text{Pr}(P | D) \geq \alpha \vee \text{Pr}(\neg P | D) \geq \alpha$$

is true for all probability models W as long as the threshold likelihood α is chosen $\leq 1/2$. Similarly,

$$\text{Tr}_D(\neg\text{LG}Q \vee \neg\text{LG}\neg Q) = \text{Pr}(Q | D) < \alpha \vee \text{Pr}(\neg Q | D) < \alpha$$

is true for all probability models as long as $\alpha > 1/2$. Thus $\text{Tr}_D(\text{LGP} \vee \text{LG}\neg\text{P} \vee \neg\text{LG}Q \vee \neg\text{LG}\neg Q)$ will be true for all choices of α . But it is easy to see that $(\text{LGP} \vee \text{LG}\neg\text{P} \vee \neg\text{LG}Q \vee \neg\text{LG}\neg Q)$ is not a valid LL sentence.

The intuitive reason behind this phenomenon is that LL can deal with situations where likelihood is interpreted as being something other than just probability. Thus, while a given LL formula may be true of any situation where L is interpreted as meaning "with probability $\geq \alpha$ ", it may not be true for some other interpretation of L . We could, for example, take LGP to mean "I have some definite information which leads me to believe that p holds with probability $\geq \alpha$ ". With this interpretation, the sentence above would not be valid.

4. Reasoning about knowledge and likelihood

We can augment LL in a straightforward way by adding modal operators for knowledge, much the same way as in [Mo, MSHI]. The syntax of the resulting language, which we call LLK, is the same as that of LL except that we add unary modal operators K_1, \dots, K_n , one for each of the "players" or "agents" $1, \dots, n$, and allow formulas of the form $K_i p$ (which is intended to mean "player i knows p "). Thus, a typical formula of LLK might be $K_i(GQ \wedge \text{LGP})$: player i knows that Q is actually the case and it is likely that P is the case.

We give semantics to LLK by extending the semantics for LL so that to each knowledge operator K_i there corresponds a binary relation \mathcal{K}_i which is reflexive, symmetric, and transitive (we remark that the assumption of symmetry gives us the axiom $\neg KP \Rightarrow K\neg KP$, and can be dropped without affecting any of the results stated below). We can think of a state and all the states reachable from it

via the \mathcal{L} relation as describing a "likelihood distribution". Two states are joined via the \mathcal{X}_i relation iff player i views them as possible likelihood distributions (rather than just possible worlds, as in [Hi,Mo,MSHI]) given his/her current knowledge. Further details, as well as proofs of the technical results stated for LLK stated in the introduction, can be found in [HMc].

5. Conclusions

We have examined the relationship between the logic LL and probability theory. We have shown that there is a precise sense in which a restricted class of probabilistic assertions about a domain can be captured by LL formulas. However, in order to correctly deal with statements of conditional probability, we must specifically list all the situations in which the conclusion may not hold. The failure to do so in informal human reasoning is frequently the cause of the nonmonotonicity so often observed in such reasoning. (However, we note here in passing a number of the problems which [McD] suggests can be dealt with by nonmonotonic logic can also be dealt with by LL, in a completely monotonic fashion. See [HR] for further discussion on this point.)

Even the restricted class of probabilistic assertions which can be dealt with by LL should be enough for many practical applications. Indeed, we view the translation from probability assertions into LL described in Section 3 as a practical tool: a discipline which forces a practitioner to list explicitly all the exceptions to his rules. Of course, this method does not guarantee correctness. If an exception is omitted, then any conclusion made using that rule may be invalid. But, whenever a conclusion is retracted, it should be possible to find the missing exception and correct the rule appropriately.

As the discussion after Theorem 1 suggests, LL seems to be able to express some notions of likelihood which probability theory cannot. This may make it applicable in contexts where probability theory is not. It would be interesting to know whether LL is able to capture other notions of reasoning about uncertainty, such as *possibility theory* ([Za1]) or belief functions ([Sh]). (See the survey paper by Prade [Pr] for a thorough discussion of various approaches to modelling reasoning about uncertainty). A number of other interesting open questions regarding Theorem 1 remain. Is there a semantics for LL, or an interpretation for L, for which a soundness and completeness result in the spirit of Theorem 1 is provable? Can we give nonstandard LL formulas a reasonable interpretation? Is there a reasonable syntax for LL in which, in some sense, all formulas are standard?

An alternative approach to reasoning about likelihood is fuzzy logic. Indeed, fuzzy logic has attempted to provide a framework for reasoning about notions such as "most", "few", "likely", and "several", which are common occurrences in natural language (cf. [Za2]). However, although the syntax of the examples in [Za2] uses these natural language notions, the semantics is still quantitative. It would be interesting to see if LL or LLK could be extended in a reasonable way to deal with the type of examples considered by Zadeh in [Za2].

Another rich area for further work is simultaneous reasoning about knowledge and likelihood. LLK provides a first step, but does not allow, for example, statements of the form "p is more likely than q". Gärdenfors ([Ga]) presents a modal logic QP where we can say "p is more likely than q", but not "p is likely". The axioms of QP seem more complicated than those of LL, and although QP is decidable,

it seems that the decision procedure would be quite complex. More research needs to be done to find an appropriate logic that is both formally and epistemologically adequate.

References

- [DBS] R. Davis, B. Buchanan, and E. Shortliffe, Production rules as a representation for a knowledge-based consultation system, *Artificial Intelligence* 8, 1977, pp. 15-45.
- [Ga] P. Gärdenfors, Qualitative probability as an intensional logic, *Journal of Philosophical Logic* 4, 1975, pp. 171-185.
- [HM] J. Y. Halpern and Y. O. Moses, Towards a theory of knowledge and ignorance, manuscript in preparation, 1984.
- [HMc] J. Y. Halpern and D. A. McAllester, Likelihood, probability, and knowledge, IBM RJ4313, 1984.
- [HR] J. Y. Halpern and M. O. Rabin, A logic to reason about likelihood, in "Proceedings of the 15th Annual Symposium on the Theory of Computing", 1983, pp. 310-319.
- [Hi] J. Hintikka, *Knowledge and Belief*, Cornell University Press, 1962.
- [Li] W. Lipski, On the logic of incomplete information, in "Proceedings of the 6th International Symposium on Mathematical Foundations of Computer Science", *Lecture Notes in Computer Science* 53, Springer-Verlag, 1977.
- [McC] J. McCarthy, Circumscription - a form on non-monotonic reasoning, *Artificial Intelligence*, 13, 1,2, 1980.
- [MH] J. McCarthy and P. Hayes, Some philosophical problems from the standpoint of artificial intelligence, in *Machine Intelligence* 4, (ed. D. Michie), American Elsevier, 1969, pp. 463-502.
- [MSHI] J. McCarthy, M. Sato, T. Hayashi, and S. Igarishi, On the model theory of knowledge, Stanford AI Laboratory, Memo AIM-312, 1978.
- [McD] D. V. McDermott, Nonmonotonic logic II: nonmonotonic modal theories, *JACM*, 29:1, 1982, pp. 33-57.
- [Mo] R. Moore, Reasoning about knowledge and action, SRI AI Center Technical Note 191, 1983.
- [Pr] H. Prade, Quantitative methods in approximate and plausible reasoning: the state of the art, Technical Report, Univ. P. Sabatier, Toulouse, 1984.
- [Re] R. Reiter, Towards a logical reconstruction of relational database theory, in *Conceptual Modelling: Perspectives from Artificial Intelligence, Databases, and Programming Languages*, (M. L. Brodie, J. Mylopoulos, and J. Schmidt, eds.), Springer-Verlag, 1984, pp. 191-233.
- [Sh] G. Shafer, *A Mathematical Theory of Evidence*, Princeton University Press, 1976.
- [SP] P. Szolovits and S. G. Pauker, Categorical and probabilistic reasoning in medical diagnosis, *Artificial Intelligence* 11, 1978, pp. 115-144.
- [Za1] L. A. Zadeh, Fuzzy sets as a basis for a theory of possibility, *Fuzzy Sets and Systems* 1, pp. 3-28, 1978.
- [Za2] L. A. Zadeh, Possibility theory and soft data analysis, in *Mathematical Frontiers of the Social and Policy Sciences* (L. M. Cobb and R. M. Thrall, eds.), A.A.A.S Selected Symposium, Vol. 54, Westview Press, Boulder, Co., 1981, pp. 69-129.