Probabilistic Reasoning under Ignorance

Marco Ramoni
Cognitive Studies in Medicine
McGill University
Montreal, Canada H3A 1A3
marco@violet.psych.mcgill.ca

Alberto Riva
Dipartimento di Informatica
Università di Pavia
I-27100 Pavia, Italy
alb@ipvaim.unipv.it

Vimla L. Patel
Cognitive Studies in Medicine
McGill University
Montreal, Canada H3A 1A3
patel@hebb.psych.mcgill.ca

Abstract

The representation of ignorance is a long standing challenge for researchers in probability and decision theory. During the past decade, Artificial Intelligence researchers have developed a class of reasoning systems, called *Truth Maintenance Systems*, which are able to reason on the basis of incomplete information. In this paper we will describe a new method for dealing with partially specified probabilistic models, by extending a *logic-based* truth maintenance method from Boolean truth-values to probability intervals. Then we will illustrate how this method can be used to represent *Bayesian Belief Networks* — one of the best known formalisms to reason under uncertainty — thus producing a new class of Bayesian Belief Networks, called *Ignorant Belief Networks*, able to reason on the basis of partially specified prior and conditional probabilities. Finally, we will discuss how this new method relates to some theoretical intuitions and empirical findings in decision theory and cognitive science.

Introduction

It is well known that classical Bayesian probability theory provides a *normative* theory of rational belief, but it fails to provide an adequate descriptive account of how human agents actually behave. This tension between its normative and its descriptive characters arises, at least in part, from the inability to distinguish between *uncertainty* and *ignorance*, since it requires to express in a single number both the belief about an event and the reliability of such a belief. Moreover, representing uncertainty in a probabilistic framework requires a large amount of information: the amount of probability estimates needed to define a complete *joint probability distribution* grows exponentially with the number of elements involved in the distribution, and this information is not always available.

During the past decade, Artificial Intelligence researchers have developed a class of reasoning systems, called *Truth Maintenance Systems* (TMSs) (Forbus and de Kleer, 1992), which are able to deal with partially specified knowledge and to perform inferences in the growth of information.

This paper describes a new method for dealing with partial information in probabilistic reasoning, by extending to intervals a *logic-based* truth maintenance method called *Boolean Constraint Propagation* (McAllester, 1990). This method will allow us to explicitly distinguish between uncertainty and ignorance. After the description of some of the theoretical and empirical motivations behind the effort of representing ignorance in probabilistic reasoning, we will outline this method and then we will show how it can be used to represent a class of probabilistic reasoning systems called *Bayesian Belief Networks* (BBNs) (Pearl, 1988), thus producing a new class of BBNs called *Ignorant Belief Networks* (IBNs), able to reason on the basis of partially specified prior and conditional probabilities. Finally, we will discuss how this new method relates to some theoretical intuitions and empirical findings in decision theory and cognitive science.

Background

A fundamental axiom of the classical Bayesian theory assumes that the beliefs of a rational agent can be expressed by a single probability distribution. The best known support for this "single distribution" assumption is the *Dutch Book Theorem*, which states that if somebody violates this assumption, he "could have a book against him by a cunning bettor, and would then stand to lose in any event" (Ramsey, 1931). This theorem provides the classical Bayesian theory with a *normative* character, by proving that a rational agent *should* behave according to its axioms, including the "single distribution" assumption.

Theoretical Arguments

Several theoretical arguments have been addressed against the assumption that a rational agent's beliefs can be expressed by a single probability distribution. A strong argument is that, under this assumption, classical Bayesian theory is unable to represent the *reliability*, or *ambiguity*, of the available information, but nonetheless this reliability affects the choices of a rational agent.

The most famous example of this effect is a well-known paradox due to Daniel Ellsberg (1961). Imagine an urn containing 30 red balls and 60 balls that can be either black or yellow, in unknown proportion. A ball is drawn at random from the urn. Ellsberg describes two choice situations, each containing two alternatives, summarized in Table 1. In the first situation, the agent is asked to choose between two alternatives: in the first alternative (I), he will receive 100 if the drawn ball is red and nothing if it is black or yellow; in the second case (II), he will receive 100 if the drawn ball is red and nothing oth-

	Red	Black	Yellow
	30	60	
I	100	0	0
II	0	100	0
III	100	0	100
IV	0	100	100

Table 1: Matrix for Ellsberg's Paradox

erwise. The second choice situation presents again two alternatives: in the first alternative (III), the agent will receive 100 if the drawn ball is red or yellow, and nothing if it is black, while in the second one (IV), he will receive 100 if the ball is black or yellow, and nothing otherwise.

Ellsberg claims that the intuitive response to these decision situations is that I is preferred over II and IV is preferred over III. This response violates the Sure-thing Principle, which requires the ordering of I to II to be preserved in III and IV, since the two pairs differ only in their third column, constant for each pair. Ellsberg suggests that this pattern of behavior is due to the *quality* of available information: the agent is sure about the proportion of the red balls in the urns, but the proportion of black balls is underconstrained (it could be anything between zero and two-thirds), thus preventing the agent to rule out a number of possible distributions.

Several authors (Good, 1962; Kyburg, 1968; Dempster, 1967; Schick, 1979) have suggested that this ambiguity or ignorance about the reliability of probabilistic information can be represented in terms of probability *intervals*. Levi (1974) interprets these intervals as sets of probability distributions. The intuition behind this theory is that even if an agent cannot choose a single probability distribution, his knowledge can constrain the set of possible distributions. Kyburg (1983) outlines a complete decision theory in which the standard probability function associated to a proposition is constrained between the functions defining the lower and upper bounds of a probability interval.

Empirical Findings

Since the seminal work by Becker and Brownson (1964), several empirical studies suggest that the "single-distribution" assumption has no descriptive value. In their work, Becker and Brownson provided the subjects with limited information about the distributions of white and black balls in an urn, giving them just *ranges* of possible distributions, and they found a kind of monetary "payoff of ignorance": the amount of money that the subjects was willing to pay to avoid ambiguity was consistently related to the amount of ambiguity of the available information.

Einhorn and Hogarth (1985) proposed an influential descriptive model of reasoning under ignorance in which an initial estimate provides the anchor and adjustments are made to cope with the missing information, and they describe several experiments supporting the basic claim of Ellsberg's Paradox. Recent results were found consistent with the predictions derived from their conceptualization of ambiguity (Hinsz and Tindale, 1992).

Results obtained by Curly and Yates (1989) and by Stasson, Hawkes, Smith and Lakey (1993) show that ignorance plays an important and autonomous role in decision processes under uncertainty, and they provide further evidence for the intuition behind Ellsberg's Paradox.

Belief Maintenance

TMSs are reasoning systems that incrementally record justifications for beliefs and propagate binary truth-values along chains of justifications. TMSs that are able to reason on the basis of probabilistic rather than binary truth-values are called *Belief Maintenance Systems* (BMSs). This section summarizes the description of a BMS based on probabilistic logic able to reason on the basis of incomplete probabilistic information. From a logical point of view, it can be regarded as an extensional system based on a set-theoretic interpretation of probability.

Logic-based Belief Maintenance

A *logic-based* BMS (LBMS) (Ramoni and Riva, 1993) is a BMS in which Boolean connectives of standard logic act as constraints on the probabilistic truth-values of propositions. The LBMS can be regarded as a generalization to interval truth-values of the Boolean Constraint Propagation method (McAllester, 1990). The LBMS manipulates two basic kinds of structures: *propositions*, $a_1, \ldots, a_n$, representing atomic propositions of a propositional language, and *clauses*, of the form $(a_1 \vee a_2 \vee \ldots)$, representing finite disjunctions of (negated or unnegated) atomic propositions.

Probabilistic logic (Nilsson, 1986) provides a semantic framework for extending the standard (Boolean) concept of satisfaction to a probabilistic one, that can be regarded as a generalization of the set-theoretic interpretation of the probability of a proposition. The probability $P(a_i)$ of a proposition a_i is bounded by the following inequality:

$$P(a_j) + P(a_j \supset a_i) - 1 \leq P(a_i) \leq P(a_j \supset a_i) \quad (1)$$

Inequality (1) may be regarded as the probabilistic interpretation of *modus ponens*. Since $(a_j \supset a_i) \equiv (\neg a_j \vee a_i)$, (2) is a special case of a more general inequality that applies to any set of propositions in disjunctive form (i.e. a clause). Let $C = \bigvee_{i=1}^{n} a_i$ be a clause. The probability of a_i is bounded by the following inequality:

$$P(C) - \sum_{j \neq i} Pa_j \leq P(a_i) \leq P(C) \quad (2)$$

The right hand side of (2) is straightforward: no proposition may have a probability greater than the probability of any disjunction it is a part of. In set-theoretic terms, this means that a set cannot be larger than its union with other sets. The left hand side states that, when the sum of the maximum probability of all the propositions in C but one does not reach the probability of the clause, the minimum probability of the remaining proposition is forced to cover the difference.

Unfortunately, the constraints directly derived from inequality (2) turn out to be too weak: the bounds they produce are too wide, thus including inconsistent values. The weakness of the constraints derived from (2) arises from too strong an enforcement of their *locality* based on the assumption that, in the set-theoretic interpretation of a clause, the intersection of all propositions is always empty. It is apparent that this assumption is too strong.

Generalizing definition (2) to interval truth-values, we derived a set of constraints on the minimum and maximum probability of propositions (Ramoni and Riva, 1993). If we define $P_*(()a_i)$ and $P^*(()a_i)$ as denoting the minimum and the maximum probability of the proposition a_i, then for each clause C in which it appears:

$$\begin{array}{ll} 1. & P_*(a_i) \geq P_*(C) + \mathcal{F}_C - \sum_{j \neq i} P^*(a_j) \\ 2. & P_*(a_i) \geq 1 - \sum_j (1 - P_*(\neg a_i \vee \psi_j)) \end{array} \quad (3)$$

Inequality (3.1) simply enforces the left hand side of (2) by dropping the assumption that all the propositions have to be pairwise disjoint and that their intersection in a set-theoretic interpretation (we call it *overlapping factor* of the clause C and denote it with $\mathcal{F}_C$) has to be empty.

Inequality (3.2) is directly derived from the well known *Additivity axiom* and states that if a_i is an atomic proposition, and $\{\phi_1, \ldots, \phi_{2^n}\}$ is the set of all the conjunctions that contain all possible combinations of the same n atomic propositions negated and unnegated, then:

$$P(a_i) = \sum_{j=1}^{2^n} P(a_i \wedge \phi_j) \quad (4)$$

The constraint (3.2) subsumes the right hand side of (2), which states that the maximum probability of each proposition in a clause C cannot be higher than the maximum probability of C. Henceforth, constraint (3.2) can be regarded as an enforcement of the inequality (2).

Propagating the constraints (3) over a network of clauses is quite easy. In the LBMS, each proposition is labeled with a set of possible values, and the constraints (in our case, the application of the above defined constraints to the clauses) are used to restrict this set. The LBMS can exhibit this behavior because if a clause is satisfied for a given truth-value of a proposition $P(a_i) = [\alpha_*, \alpha^*]$, it will be satisfied for any subset of $[\alpha_*, \alpha^*]$. This property, which is implicit in the form of the inequalities in constraints (3), implies a monotonic narrowing of the intervals, thus ensuring the incrementality of the LBMS.

The most important feature of the LBMS is the ability to reason from any subset of the set of clauses representing a joint probability distribution, by bounding the probability of the propositions within probability intervals, and incrementally narrowing these intervals as more information becomes available. The assignment of *probability bounds* to propositions and clauses expresses the degree of ignorance about them.

Representing Conditionals

A common criticism addressed to extensional systems based on probabilistic logic is their difficulty in representing assignments of *conditional* rather than *absolute* probabilities. Pearl (1988) points out that a more intuitive representation of the *modus ponens* in a probabilistic framework is given by the specification of $P(a_j)$ and $P(a_i|a_j)$ rather than by implication $P(a_j \supset a_i)$, because $P(a_j \supset a_i)$ does not properly capture the meaning of the linguistic statement "if a_j then a_i". For instance, if we want to express that some rare event a_j has a likely consequence a_i and we state $P(a_j) = [.01\ .01]$ and $P(a_j \supset a_i) = [.9\ .9]$, we find that the two sentences are inconsistent. If we represent *modus ponens* using the conditional probability $P(a_i|a_j)$ rather than implication $P(a_j \supset a_i)$, the probability of a_i is bound by the following inequality:

$$P(a_j) \cdot P(a_i|a_j) \leq P(a_i) \leq 1 \quad (5)$$

In the previous example, if we write $P(a_j) = [.01\ .01]$ and $P(a_i|a_j) = [.9\ .9]$, inequality (5) produces the bound $P(a_i) = [.09\ 1]$, which is more intuitive. Modus ponens is, in nature, a metalinguistic rule of inference. Henceforth, following the approach of Kyburg (1983), we will represent conditionals as metalinguistic derivation rules, rather than as connectives in the object language.

The representation of conditional probabilities in the LBMS is straightforward using the Chain Rule:

$$P(a_j) \cdot P(a_i|a_j) = P(a_i \wedge a_j) \quad (6)$$

The resulting conjunction is converted in clausal form through De Morgan's laws and it is then communicated to the LBMS. For instance, the probabilistic model defined by the two conditionals $P(a_2|a_1) = [.2\ .2]$ and $P(a_2|\neg a_1) = [.6\ .6]$ with $P(a_1) = [.5\ .5]$ may be expressed by the set of clauses: $P(a_1 \vee a_2) = [.8\ .8]$, $P(a_1 \vee \neg a_2) = [.7\ .7]$, $P(\neg a_1 \vee a_2) = [.6\ .6]$, $P(\neg a_1 \vee \neg a_2) = [.9\ .9]$.

Representing Ignorance

A BBN is a direct acyclic graph in which nodes represent stochastic variables and links represent causal relationships among those variables. Each link is defined by the set of all conditional probabilities relating the parent variables (the "cause" variables) to children variables (the "effect" variables). We can identify two different kinds of ignorance that can be represented in this framework: complete ignorance about a conditional probability and partial ignorance about a conditional or prior probability in the network.

Ignorant Belief Networks

From the theory of the TMSs, the LBMS inherits the concept of *consumer*. A consumer is a forward-chained procedure attached to each proposition, that is fired when the truth-value of the proposition is changed. The BMSs theory extends the definition of consumers from Boolean to probabilistic truth-values. Using consumers, it is possible to develop a new class of BBNs based on the LBMS and henceforth able to reason with partially specified

	Conditional	P
1	[tuberculosis:yes] | [asia:no]	[.01 .01]
2*	[tuberculosis:yes] | [asia:yes]	[.05 .05]
3	[bronchitis:yes] | [smoker:yes]	[.6 .6]
4*	[bronchitis:yes] | [smoker:no]	[.3 .3]
5	[lung-cancer:yes] | [smoker:yes]	[.1 .1]
6*	[lung-cancer:yes] | [smoker:no]	[.1 .1]
7	[dyspnoea:yes] | [bronchitis:yes] ∧ [tuber-or-cancer:yes]	[.9 .9]
8	[dyspnoea:yes] | [bronchitis:yes] ∧ [tuber-or-cancer:no]	[.7 .7]
9	[dyspnoea:yes] | [bronchitis:no] ∧ [tuber-or-cancer:yes]	[.8 .8]
10*	[dyspnoea:yes] | [bronchitis:no] ∧ [tuber-or-cancer:no]	[.1 .1]
11	[tuber-or-cancer:yes] | [tuberculosis:yes] ∧ [lung-cancer:yes]	[1 1]
12	[tuber-or-cancer:yes] | [tuberculosis:yes] ∧ [lung-cancer:no]	[1 1]
13	[tuber-or-cancer:yes] | [tuberculosis:no] ∧ [lung-cancer:yes]	[1 1]
14	[tuber-or-cancer:yes] | [tuberculosis:no] ∧ [lung-cancer:no]	[0 0]
15	[x-ray:yes] | [tuber-or-cancer:yes]	[.98 .98]
16*	[x-ray:yes] | [tuber-or-cancer:no]	[.05 .05]

Table 2: Conditional probabilities defining the network of the example.

causal links (i.e. lacking some conditional probabilities) and interval probability values. We call these BBNs *Ignorant Belief Networks*. In this framework, IBNs act as a high-level knowledge representation language, while the computation and the propagation of probabilities are performed by the LBMS.

In a BBN, each variable is defined by a set of *states* representing the assignment of a value to the variable. Each *state* is evaluated by a probability value. All the states of a variable are mutually exclusive and exhaustive: the probability values assigned to all the states in a variable have to sum to unit. In an IBN, when a variable is defined, each state is communicated to the LBMS as a proposition and a set of clauses and consumers is installed to ensure that its states are mutually exclusive and exhaustive.

Causal relations are defined by conditional probabilities among states. In an IBN, a conditional $P(e|Cx)$ is represented as a consumer attached to each proposition representing a state in the context Cx. When the probability value of all states in Cx is assigned, the consumer communicates to the LBMS the two different clauses resulting from the application of the De Morgan's laws to $(C \wedge e)$ and $(C \wedge \neg e)$. $P(C \wedge e)$ and $P(C \wedge \neg e)$ are calculated according to a version of the Chain Rule extended to intervals:

$$P_*(C \wedge e) = P_*(Cx) \cdot P_*(e|Cx)$$
$$P^*(C \wedge e) = P^*(Cx) \cdot P^*(e|Cx)$$
$$P_*(C \wedge \neg e) = P_*(Cx) \cdot (1 - P^*(e|Cx))$$
$$P^*(C \wedge \neg e) = P^*(Cx) \cdot (1 - P_*(e|Cx))$$

It is worth noting that since the probability of both propositions and clauses in the LBMS is represented by probability intervals, IBNs are endowed with the ability to express both interval conditional probabilities and interval prior probabilities about states. Moreover, since conditionals are locally defined and propagated, the reasoning process can start even without the full definition of the joint probability distribution. These features enable IBNs to represent both the complete ignorance of a conditional probability and the partial ignorance of a conditional or a prior probability.

An Example

In order to illustrate the functionality of IBNs, we will use the well known example depicted in Figures 1-4. The pop-up windows over the variables show, in graphical terms, the probability interval (subset of [0 1]) associated with each of their states. In each bar, the area between 0 and $P_*(.)$ is black, the area between $P^*(.)$ and 1 is white, and the area between $P_*(.)$ and $P^*(.)$ is gray. In a standard BBN, all conditional and prior probabilities reported in Table 2 are needed before any reasoning process can start. Figure 1 shows an IBN in which those conditional probabilities that are denoted by an asterisk in Table 2 are missing, and the prior probabilities of root nodes are intervals rather than point-valued: $P(\texttt{[asia : yes]}) = [.05\ .15]$ and $P(\texttt{[smoker : yes]}) = [.85\ .95]$.

Figure 2 shows a portion of the LBMS network generated by the propagation. Rectangles represent propositions and ovals are clauses. A solid arc linking a proposition to a clause means that the proposition appears unnegated in the clause, while a dashed arc means that it appears negated. The side bars display the minimum and maximum probability. The thicker border of the proposition [asia:yes] indicates that it is an assumption. The clauses $P(\texttt{[asia : yes]} \vee \texttt{[asia : no]}) = [1\ 1]$ and $P(\neg\texttt{[asia : yes]} \vee \neg\texttt{[asia : no]}) = [1\ 1]$ and the clauses $P(\texttt{[tuberculosis : yes]} \vee \texttt{[tuberculosis : no]}) = [1\ 1]$ and $P(\neg\texttt{[tuberculosis : yes]} \vee \neg\texttt{[tuberculosis : no]}) = [1\ 1]$ enforce the mutual exhaustivity and exclusivity between the propositions representing the states of the variables **Asia** and **Tuberculosis**, respectively.

The clauses $P(\texttt{[asia : yes]} \vee \neg\texttt{[asia : no]} \vee \texttt{[tuberculosis : yes]}) = [.0595\ .1585]$ and $P(\texttt{[asia : yes]} \vee \neg\texttt{[asia : no]} \vee \neg\texttt{[tuberculosis : yes]})$

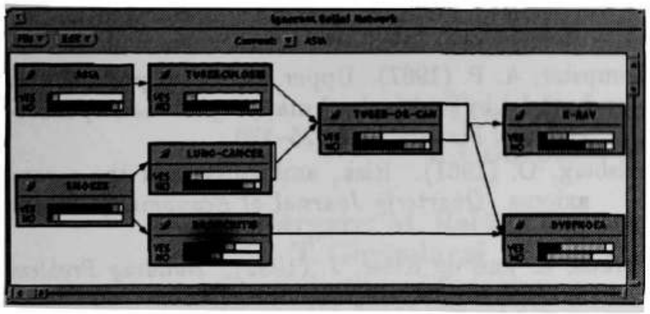

Figure 1: The IBN defined by conditionals without the asterisk in Table 2 and interval prior probabilities.

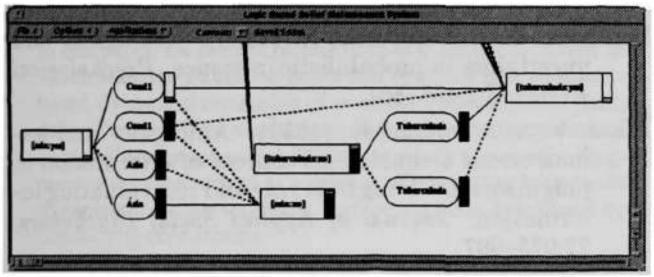

Figure 2: A section of the LBMS network defined by the propagation of consumers for the IBN in Figure 1.

= [.9905 .9915], named **Cond1**, are generated by the application of the Chain Rule with P([**asia** : **yes**]) = [.05 .15] and P([**tuberculosis** : **yes**]|[**asia** : **no**]) = [.01 .01]. Note the absence of any clause representing the application of the conditional 2 in Table 2. The underspecified probability of the proposition **[tuberculosis:yes]** is due both to the interval-valued probability of the proposition **[asia:yes]** and to the absence of the conditional P([**tuberculosis** : **yes**]|[**asia** : **yes**]) = [.05 .05].

As a matter of fact, Figure 3 shows that when point-valued prior probabilities for the states of root variables (P([**asia** : **yes**]) =[.1 .1] and P([**smoker** : **yes**]) =[.9 .9]) are assumed, the probability of **[tuberculosis:yes]** remains underspecified. Nonetheless, Figure 3 shows the monotonic narrowing of all probability intervals in the network: due to the incremental character of the LBMS, all intervals in Figure 3 are subset of intervals depicted in Figure 1. Finally, when the missing conditional probabilities of Table 2 are communicated to the IBN, the intervals degenerate to point-valued probabilities, and the IBN converges to the values of a standard BBN, as depicted in Figure 4.

Discussion

We have introduced a new method for dealing with partially specified probabilistic models in intelligent systems and we have applied it to develop a new class of BBNs able to reason on the basis of an explicit representation of ignorance.

We have applied the IBNs to forecasting blood glucose concentration in insulin-dependent diabetic patients

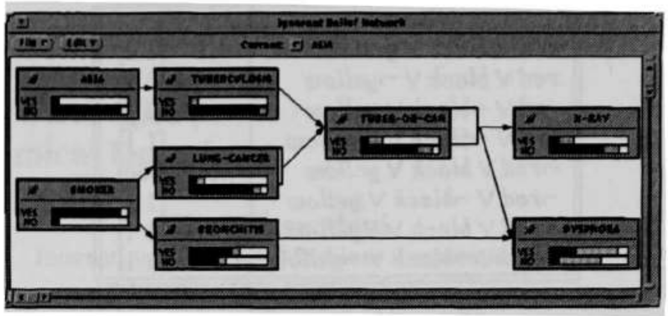

Figure 3: The IBN defined by conditionals without the asterisk in Table 2 and point-valued prior probabilities.

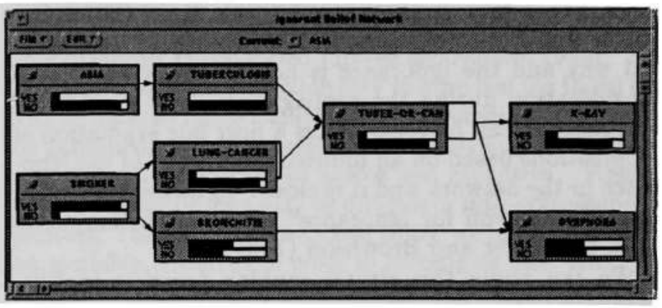

Figure 4: The IBN defined by all conditionals in Table 2 and point-valued prior probabilities.

using underspecified probabilistic models directly derived from a database containing the daily follow-up of 70 insulin-dependent diabetic patients, in which a very small subset of the complete conditional model needed to define a BBN was available. Instead of the 19200 conditional probability required, only 2262 were available (that is, less than 12%), and most of them were affected by ignorance (the mean difference between the maximum and minimum probability of the conditionals was .19) (Ramoni *et al.*, 1994). Still, the system was able to reason and to predict blood glucose values, taking into account the ignorance about the distributions.

The crucial problem we found in the representation of ignorance was the discrimination between two states having intersecting probability intervals. This is a major question for interval-based approaches to decision theory. Kyburg (1983) introduces a general rule to for decision making under ignorance: "It is rational to reject any choice for which there exists another choice whose minimum expected utility exceeds its own maximum expected utility." If we assume the utility function as constant, we can choose as "predicted" a state a_i if there is no alternative state a_j whose maximum probability exceeds the minimum probability of a_i. Unfortunately, this is a very rare case, and the rule proposed by Kyburg is unable to legislate when the intervals are not disjoint.

However, interpreting the width of the interval as a measure of ignorance, we can rank the states of each variable according to a score that is proportional to their mean probability and inversely proportional to the igno-

Clause	P
$red \vee black \vee yellow$	[1 1]
$red \vee black \vee \neg yellow$	[.333 1]
$red \vee \neg black \vee yellow$	[.333 1]
$red \vee \neg black \vee \neg yellow$	[1 1]
$\neg red \vee black \vee yellow$	[.666 .666]
$\neg red \vee \neg black \vee yellow$	[1 1]
$\neg red \vee black \vee \neg yellow$	[1 1]
$\neg red \vee \neg black \vee \neg yellow$	[1 1]

Table 3: Clausal representation of the probability distributions for Ellsberg's Paradox.

rance about their probability. On this view, the probability values are propagated by the LBMS in a categorical way and the ignorance is not explicitly taken into consideration during the propagation. The discrimination method can be regarded as a *post hoc* evaluation of propositions based on an induced estimation of the ignorance in the network and it is closed, in its nature, to the idea of a "payoff for ignorance" found in the empirical work by Becker and Brownson (1964).

We can apply this simple ranking function to Ellsberg's Paradox, since the assumed monetary payoffs are constant. Table 3 shows the clausal representation of the probability distributions underling the paradox. The LBMS calculates the following probability intervals for the propositions *red*, *black*, and *yellow*: $P(red) = [.333\ .333]$, $P(black) = [0\ .666]$, and $P(yellow) = [0\ .666]$.

The mean values of the intervals associated to the statements representing the alternatives in Ellsberg's Paradox are all equal, and just the width of the intervals, defining the ignorance about their probability, is different. It is easy to see that, following our simple ranking rule, even the smallest penalty given to the statements affected by ignorance will lead to prefer the alternative I over the alternative II and the alternative IV over the alternative III in Table 1, in agreement with the intuition behind Ellsberg's Paradox and the empirical findings supporting it.

Acknowledgments

This research was supported in part by a grant from the Social Science and Humanities Council of Canada (410 92 1535) to Vimla L. Patel, and by the AIM Programme of the Commission of the European Communities (A2034). The authors thank Riccardo Bellazzi, Mario Stefanelli, and Tony Marley for their helpful suggestions and discussions. The use of the Sun Sparc 10 Workstation was generously provided by Dr. Marley.

References

Becker, S. and Brownson, F. (1964). What price ambiguity? or the role of ambiguity in decision making. *Journal of Political Economy*, 74:801–812.

Curley, S. and Yates, J. (1989). An empirical evaluation of descriptive models of ambiguity reactions in choice situations. *Journal of Mathematical Psychology*, 33:397–427.

Dempster, A. P. (1967). Upper and lower probabilities induced by multivalued mapping. *Annals of Mathematical Statistics*, 38:325–339.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics*, 75:643–669.

Forbus, K. and de Kleer, J. (1992). *Building Problem Solvers.* MIT Press, Cambridge, MA.

Good, I. (1962). Subjective probability as a measure of a non-measurable set. In Nagel, E., Suppes, P., and Tarsky, A., editors, *Logic, Methodology and Philosophy of Science*, pages 319–329. Stanford University Press, Stanford, CA.

Heinhorn, H. and Hogarth, R. (1985). Ambiguity and uncertainty in probabilistic inference. *Psychological Review*, 92:433–461.

Hinsz, V. and Tindale, R. (1992). Ambiguity and human versus technological sources of information in judgments involving base rate and individuating information. *Journal of Applied Social Psychology*, 22:973–997.

Kyburg, H. (1968). Bets and beliefs. *American Philosophical Quarterly*, 5:63–78.

Kyburg, H. (1983). Rational belief. *Behavioral and Brain Sciences*, 6:231–273.

Levi, I. (1974). On indeterminate probabilities. *Journal of Philosophy*, 71:391–418.

McAllester, D. (1990). Truth maintenance. In *Proceedings of the 8th AAAI*, pages 1109–1115.

Nilsson, N. (1986). Probabilistic logic. *Artificial Intelligence*, 28:71–87.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of plausible inference.* Morgan Kaufmann, San Mateo, CA.

Ramoni, M. and Riva, A. (1993). Belief maintenance with probabilistic logic. In *Proceedings of the AAAI Fall Symposium on Automated Deduction in Non Standard Logics*, Raleigh, NC. AAAI.

Ramoni, M., Riva, A., Stefanelli, M., and Patel, V. (1994). Forecasting glucose concentration in diabetic patients using ignorant belief networks. In *Proceedings of the AAAI Spring Symposium on Artificial Intelligence in Medicine*, Stanford, CA. AAAI.

Ramsey, F. (1931). Truth and probability. In Braithwaite, R., editor, *The Foundations of Mathematics and Other Logical Essays*, pages 156–198. Routledge and Kegan Paul, London.

Schick, F. (1979). Self-knowledge, uncertainty, and choice. *British Journal for the Philosophy of Science*, 30:235–252.

Stasson, M., Hawkes, W., Smith, H., and Lakey, W. (1993). The effects of probability ambiguity on preferences for uncertain two-outcome prospects. *Bulletin of the Psychonomic Society*, 31:624–626.