The Knowledge Engineering Review, Vol. 11:1, 1996, 73–83

Book reviews

The uncertain reasoner's companion: a mathematical perspective by Jeff B. Paris, Cambridge University Press, Cambridge, UK, 1994, pp 212. Volume 39 in Cambridge Tracts in Theoretical Computer Science.

In the past 15 years, there has been a very active field of research devoted to probabilistic reasoning in artificial intelligence. The most visible achievement of probabilistic reasoning is the Bayesian network that has, to some extent, put the certainty factor approach to uncertainty handling in expert systems on a sound basis. However there is more to probabilistic reasoning than Bayesian networks, because the latter rely on very restrictive assumptions when construed as a representation of expert knowledge. The book by Jeff Paris tries to bridge the gap between expert knowledge and its apparently remote, artificial representation as a probability distribution on a set of interpretations of a logical language. This is a very personal book because the author develops his own view of probabilistic reasoning and proposes a refreshing analysis of principles that should guide probabilistic knowledge representation and inference. The validity of these principles extends beyond the field of probability theory.

Analysis

This rather short but very rich monograph is made of 12 chapters. It starts with a motivating situation in terms of a knowledge base that contains statements in natural language that a medical doctor typically could provide when describing domain knowledge. This is a real problem of question answering in a medical information system. The basic assumption is that uncertain statements translate into linear constraints on conditional and unconditional measures of uncertainty, defined on a logical, propositional language. These constraints delimit a set of uncertainty functions defined on this propositional language. These constraints delimit a set of uncertainty functions defined on this propositional language. Each of these uncertainty functions is said to be compatible with the knowledge base. This view leads to a mathematically rigorous model of the knowledge base, that is general enough to encompass several uncertainty theories.

On this basis, the book proceeds in two directions: on the one hand, it presents justifications for several uncertainty calculi with a strong preference in favor of probability but avoiding Bayesian dogmatism. On the other hand, the author claims that common sense meta-knowledge exists that enables a single uncertainty measure compatible with the knowledge base to be selected, for the purpose of question-answering. There are four chapters two of which are devoted to justifications of probability theory, the two others dealing with belief functions and what the author calls "truth-functional belief". The author then focuses almost exclusively on purely probabilistic representations. Chapters 6 and 7 are the core of the book because they borrow a lot from the author's own concerns and results. They present several techniques for selecting a single probability measure as representative of a set of probabilities (induced by the knowledge base) and several principles that a selection process should satisfy. The Paris principles for selecting a probability distribution out of a convex family of such distributons (as induced by a probabilistic knowledge base K) are as follows:

- Equivalence principle: two knowledge bases yielding the same probability set should produce the same selection;

- Irrelevant information: if K_1 and K_2 are knowledge bases whose languages are disjoint then the selected probability of a formula that is expressible in the language of K_1 should not change when adding K_2 K_1;
- Continuity: any microscopic change of numbers appearing in K_1 should lead to miscroscopic change in the value of the selected probability of any formula;
- Open-mindedness principle: if it is consistent with K that the probability of a formula be positive, then the selected probability of this formula should be positive as well;
- Obstinacy principle: if the probability selected from K_1 is in the probability set of K_2, then the same probability should be selected from $K_1 \cup K_2$;
- Relativisation: conditioning the selected probability on an event whose probability is stated as positive in the knowledge base should not be affected by adding to this knowledge base statements conditioned on the negation of this event;
- Independence: the selected probability should assume conditional independence between propositions where possible;
- Atomicity: changing a propositional variable into a formula by augmenting the language should not affect the conditional probability.

The authors shows that the maximal entropy selection is the only one that satisfies all principles, but atomicity. The as pointed out in the book.

On such a basis, the author then tries to address the question-answering problem via Bayes rule of conditioning applied to the selected uncertainty measure. The case when the input information (pertaining to the patient state, in the medical example) is uncertain is tackled by Jeffrey's rule of updating whose justification via minimising cross-entropy is recalled. The topic of Bayesian networks is only touched upon in Chapter 9, as a by-product of adding independence assumptions to a probabilistic knowledge-base. Adding independence assumptions is another way of selecting a relevant probability measure. It is recalled that adding too-many independence assumptions trivialises the potential inferences, while Bayesian nets add just the right amount of such assumptions so as to select a single, but general enough, probability distribution compatible with a knowledge base.

Very logically, the next step in the analysis is, once the problem of question answering in an uncertain knowledge based has been stated, and its solutions mathematically well-defined, to see if the ensuing calculations are tractable. Chapter 10 proves that, generally, they are not. This chapter provides several deep results on worst case computation complexity evaluations for problems such as consistency testing of an uncertain knowledge base, selecting a probability measure and computing its value for a query, and question answering on a Bayesian network. Overall these results sound very pessimistic. The two last chapters are devoted to the extension of probabilistic knowledge bases to the predicate calculus, and presents the ideas and results of Carnap. The aim of these chapters is to propose this extension mainly as an open problem since effective results are given only for very limited types of knowledge bases (especially empty ones).

Critique

This book appears as a useful complement to the probabilistic AI literature which so far has been almost exclusively on the Bayesian side. The Uncertain Reasoner's Companion is not so since it acknowledges the idea that a set of probability measures, induced by constraints, is a more natural representation of an uncertain knowledge base than a unique distribution. It advocates the use of a single distribution, but it does it in a much more convincing way than usual because it provides an in-depth discussion about principles that can lead us to select this distribution.

This book is also a mathematical treatise that tries to be as self-contained as possible. Any non-mathematically sophisticated reader may find it difficult to read some chapters although the tutorial merits of the book are unusually high. It especially puts together material that has sometimes acquired a mythical value in the field, things that many people talk about, often without really dealing with them in a precise way. For instance, Chapter 3 contains a full proof of Cox's theorem, that has been used by probabilists as a proof of the non-existence of alternative uncertainty calculi. Namely, from simple axioms that define conditioning, plus the assumption that the belief in a proposition is a function of the belief in the opposite proposition, and a density assumption on top, this theorem establishes that probability theory is the only mathematically coherent uncertainty approach. Paris points out that many authors forget about the density assumption. Chapter 4 contains a full proof that Bayesian conditioning applied to a belief function (viewed as a lower probability) yields a belief function, a result independently established by Fagin and Halpern and Jaffray. Chapter 5 is a full mathematical introduction to fussy set connectives (triangular norms and conorms). This effort towards self-containedness recommends this book to any mathematically oriented scholar in the field of uncertain reasoning.

The most original part of the monograph, which can be viewed as a major contribution to the literature is Chapter 7, about principles of uncertain reasoning. The author first casts probabilistic constraint propagation into a logical setting and claims that the commonsense principles recalled here above, added to domain knowledge lead to the selection of a unique probability distribution. According to Paris, a conclusion $P(\theta) = a$ can be derived from a probabilistic knowledge base K if the "best" probability distribution compatible with K is such that $P(\theta) = a$. One of Paris' achievements is to show that maximum entropy receives new justification from the commonsense principles of probability selection, as opposed to other apparently natural selection methods such as the centre of mass.

The centre of mass process is very much criticised in Paris' book because it is not language-independent, and violates some other principles of commonsense selection. However, the counterexample of the Brandeis dice problem raises an interesting debate. It consists of a three-sided die, numbered 1, 2, 3, whose expected score is 2. The question is to guess the probability of the die landing 3 on the next throw. An expected throw of 2 means that 1 and 3 are equally likely. If x_1, x_2, x_3 are the probabilities of landing 1, 2 and 3, respectively, then the equations linking x_1, x_2 and x_3 are $x_1 = x_3$ and $x_1 + x_2 + x_3 = 1$. Letting $x_1 = a$, we get $x_1 = x_3 = a$ and $x_2 = 1 - 2a$, it is easy to check that the probability selected by the centre of mass procedure gives $x_1 = x_3 = 1/4$ while $x_2 = 1/2$. This result is paradoxical because "everything we know" (here, an expected throw of 2) "is consistent with a die being fair" (p. 73).

This line of reasoning is convincing only if the probability bounds induced by the constraints are not viewed as degrees of uncertainty. Indeed, suppose they are, namely, suppose lower bounds represent degrees of belief and upper bounds represent degrees of plausibility. Then notice that the lower bounds $P_*(1), P_*(3) = P_*(2) = 0$. On the contrary the upper bounds are $P^*(1) = P^*(3) = 1/2$ and $P^*(2) = 1$. Under this view, the centre of mass inference appears more natural, since $x_2 = 1/2 > x_1 = 1/4$ reflects the fact that the plausibility of 2 ($Pl(2) = P^*(2)$) is higher than the plausibility of 1 ($Pl(1) = P^*(1)$). What this discussion points out is the distinction between upper and lower probabilities as incomplete probabilistic knowledge (where the centre of mass transformation is paradoxical), and degrees of plausibility and belief *modelled by* upper and lower probability bounds (where the centre of mass is natural). In the first situation there is no probability measure P such that $P(1) = P(3) = 1/2$ and $P(2) = 1$, so that such degrees of belief are inconsistent with probability theory. In the second situation, $Pl(1) = Pl(3) = 1/2$, $Pl(2) = 1$ is a fully acceptable statement expressing less possibility in landing 1, or landing 3, than in landing 2. Note that in belief function theory after Smets (the transferable belief model), there is a definite rejection of the incomplete probabilistic knowledge interpretation of belief functions (although this interpretation is mathematically feasible) and this rejection is motivated by the assumption that Shafer/Smets belief functions are supposed to represent (precise) beliefs. However, when Paris discusses belief functions in Chapter 4, he only talks about probability bounds thus missing the point made by the transferable belief model.

It is patent that all principles for the selection of a best probability distribution, but independence and relativisation, are not bound to the choice of an uncertainty theory. Relativisation becomes general if we modify the concept of conditioning as stated in probability, so as to stick to the counterpart to conditioning in the alternative uncertainty theory; the same applies to independence which is a typically probabilistic concept and on which little is known outside probability. Provided such an adaptation, the same principles applied to possibility theory justify the choice of a minimally specific possibility distribution from a possibilistic knowledge base, as recently proved by a student of Paris (Maung, 1993). But this possibilistic selection obeys atomicity as well. This view is strikingly similar to the one advocated in non-monotonic reasoning, especially the approaches based on model minimization (Shoham, 1988). In non-monotonic reasoning a formula θ is a consequence of knowledge base K if and only if the "best" models of K satisfy θ. In fact, common sense inference based on Lehmann's rational closure (Lehmann and Magidor, 1992) or based on infinitestimal probabilities (Pearl, 1990) also adopt the view of selecting a least committed preference structure from a family of constraints induced by default rules, and this kind of inference is fully captured by the above defined selection of a possibility distribution characterized by Maung (see Benferhat et al., 1992). Thus, there is a strong similarity between non-monotonic reasoning and Paris' approach to probabilistic reasoning.

The same methodology consisting of using constraints to represent knowledge so as to capture partial ignorance in a faithful way, and adding commonsense principles so as to select a more informative (if somewhat adventurous) belief measure for exploiting the available knowledge is also at work in belief function theory as developed by Smets (1988, 1990a,b, 1993, 1994). Smets claims that belief functions are a faithfully representation of belief at the credal level. Now for the purpose of decision, he gives commonsense principles that select a single probability distribution which can reasonably used to compute expected utility. This transformation, dubbed "pignistic" by Smets (1990), is called the centre of mass process by Paris in his book (viewing a belief function as a lower probability). This methodology is again at work in fuzzy controllers, where a fuzzy inference step is followed by a defuzzification step (e.g. Driankov et al., 1993) that selects a value for the control parameter. It is not the least merit of the book to present this two-step inference methodology with full generality, even if it is exploited only on probabilistic representations.

These remarks lead us to comment on the treatment of non-probabilistic calculi in Paris book. Non-classical uncertainty plays a minor role in the book but there are two chapters devoted to them. The presentation of these theories is not fully satisfactory for several reasons. First belief functions are treated as upper and lower probabilities. Moreover, although the two existing conditioning rules are presented with full mathematical details, no hint is given on their respective role for the purpose of inference. Actually the work of Smets (1988, 1990a,b, 1993, 1994), which is the most advanced to-date on the topic of justifying the belief function approach, is totally omitted. Similarly, possibility theory is only viewed as a special case of belief function theory and got ridden of in half a page in the chapter on truth-functional belief, while the author acknowledges the fact that possibility is not truth-functional.

As for the chapter on truth-functional belief, there is, to this reader's opinion a major flaw in it. Commenting on the truth-functionality assumption, the author says that "we are, regrettably, unaware of any satisfactory justification of this except possibly on pragmatic grounds" (p. 53). However, the author apparently never acknowledges that truth-functional belief is simply impossible because it would mean some form of homomorphism between a Boolean algebra and a subset of the unit interval with more than two elements in it. Proof of such impossibility has been established elsewhere (Dubois and Prade, 1988, 1994). Failing to acknolwedge the impossibility of truth-functional belief leads either to an unfortunate confusion between belief calculi and fuzzy logic (a confusion indeed made by Paris in this chapter since he says that a particular case of truth-functional belief calculus "is often referred to as fuzzy logic" (p. 52)), or to the rejection of fuzzy logic as mathematically inconsistent (a mistake made by Elkan (1994) for instance). However, the distinction between degrees of truth and degrees of belief was made very early by De Finetti (1936) who pointed out that degrees of belief are on the meta-level with respect to degrees of truth. So the chapter on truth-functional belief really deals with the pair $\{0,1\}$ of truth-values, and not the unit interval as the author would like us to believe, since under this chapter's assumptions, belief values collapse down to binary truth-values.

A second debatable stand taken by Paris concerns what he calls "belief revision". Paris explicitly considers the following situation: a knowledge base contains medical generic knowledge. The next patient arrives, and the medical doctor acquires particular information about this patient. A first assumption made by Paris is that this information forms a new knowledge base which is added to the generic one, in order to make inference. A second assumption is that when the particular knowledge is not uncertain and takes the form of a standard proposition to focus upon, Paris suggests that the question-answering problem can be solved by conditioning on this proposition the probability distribution selected from the generic knowledge base. Although I trust the author's competence for putting together a mathematically coherent explanation and implementation of these proposals, I have some doubts that these methods solve the question of inference in the face of new evidence, or even the one of belief revision.

First the author does not seem to acknowledge the distinction between question-answering and belief revision, perhaps because in the face of a single probability distribution representing a generic knowledge, both problems are often dealt with by means of Bayes rule of conditioning. There is a classical debate in the probabilistic literature between belief revision people (the probability kinematics school, some results of which are presented in Chapter 8) and those who view Bayesian conditioning as a change of reference class (De Finetti, noticeably). However, in the face of a set of probability measures, which is the basic representation tool in the *Uncertain Reasoner's Companion*, the use of conditioning becomes more tricky. Reading Chapter 6, one gets the impression that Paris solves a question answering problem (is the patient, on which I have some evidence, ill?) by means of a belief revision approach. Adding the particular knowledge about the patient to the generic knowledge in the knowledge base may lead to a revision of the latter that is, altering its contents. Namely, let K be a set of probabilistic constraints delimiting a probability set $\mathcal{P}_K$. If the next patient lies in reference class A (pointed out by evidence), it sounds strange to add the constraint $\mathcal{P}(A) = 1$ to K, because $P(A) = 1$ really means that *all* patients are in class A. Similarly Jeffrey's rule consists in enforcing a constraint of the form $P(A) = a \neq 0,1$, and also results in revising the generic knowledge. Maximisation of cross entropy again amounts to a revision, and, as does Jeffrey's rule, accounts for the case when the new *generic* knowledge is inconsistent with the *a priori* generic knowledge K. On the contrary the question answering problem with generic knowledge K and reference class A of which the patient is supposed to be a typical member, can be addressed by considering the set $\mathcal{P}_{K\mid A} = \{P(\cdot \mid A), P \in \mathcal{P}_K\}$ of conditional probabilities pertaining to class A. Clearly, the set $\mathcal{P}_{K\mid A}$ devoted to question-answering clearly differs from the set $(\mathcal{P}_K)^*_A = \{P \in \mathcal{P}_K, P(A) = 1\}$ whereby $\mathcal{P}_K$ is revised upon the arrival of a new constraint $P(A) = 1$. For instance, the latter can be empty (if $\sup\{P(A), P \in \mathcal{P}_K\} < 1$) when the former is not (if $\inf\{P(A), P \in \mathcal{P}_K\} > 0$).

How should the selection process act in the case where the distinction between question answering and revision is acknowledged? Rather than conditioning the previously selected probability function, it sounds more natural to operate a selection on the focused probability set, or on the revised one. If $N(\mathcal{P})$ denotes the typical selected probability from $\mathcal{P}$, then, the common sense solving the question-answering problem "does the patient belong to class C?" should be $N(\mathcal{P}_{K\mid A})(C)$ rather than $N(\mathcal{P}_K)(C \mid A)$. It might be that $N(\mathcal{P}_{K\mid A})(C)$ equates $N(\mathcal{P}_K)(C \mid A)$ when N is the maximal entropy selection. Indeed, this equality may sound like an alternative proposal to Paris' relativisation principle, and it would simplify calculations. As for revision, it is very unlikely that $N(\mathcal{P}_K)(C \mid A)$ be in accordance with the selection operated from the revised belief set. Such a failure has been observed by Smets (1990) with his pignistic transformation. These comments have no ambition to fully address these questions, but suggest that in Chapter 6, if the presented techniques are mathematically sound, they may appear semantically misleading.

Despite the above reservations, the *Uncertain Reasoner's Companion* is a valuable contribution to the literature of uncertainty in artificial intelligence. As a mathematical treatise on probabilistic reasoning it is exceptionally well-written and can serve as a textbook. It also makes useful, hard to find part of the uncertainty literature accessible in the same spot. As a reflexion on the meaning of commonsense inference, it may become a major contribution due to the systematic, axiomatic, mathematically clean, and potentially generic attempt at characterising this implicit part of our knowledge that makes us jump to plausible conclusions.

Reviewed by Didier Dubois, IRIT-CNRS, Toulouse, France

References

Benferhat, S, Dubois D and Prade, H, 1992. "Representing default rules in possibilistic logic" In: *Proc. of the 3rd Inter. Conf. on Principles of knowledge Representation and Reasoning* (*KR'92*), 673–684, Cambridge, MA, October 26–29.

De Finetti, B, 1936. "La logique de la probabilite" *Actes du Congrès Inter. de Philosophie Scientifique*, Paris. (Hermann et Cie Editions, 1936, IV1–IV9).

Driankov, D, Hellendoorn, H and Reinfrank, M, 1995. *An Introduction to Fuzzy Control,* Springer-Verlag.

Dubois, D and Prade, H, 1988. "An introduction to possibilistic and fuzzy logics" In: *Non-Standard Logics for Automated Reasoning* (P Smets, A Mamdani, D Dubois and H Prade, editors), 287–315, Academic Press.

Dubois, D and Prade, H, 1994. "Can we enforce full compositionality in uncertainty calculi?" In: *Proc. 12th US National Conf. On Artificial Intelligence* (*AAAI94*), 149–154, Seattle, WA.

Elkan, C, 1994. "The paradoxical success of fuzzy logic" *IEEE Expert* August, 3–8.

Lehmann, D and Magidor. M, 1992. "What does a conditional knowledge base entail?" *Artificial Intelligence* **55** (1) 1–60.

Maung, I, 1995. "Two characterizations of a minimum-information principle in possibilistic reasoning" *Int. J. of Approximate Reasoning* **12** 133–156.

Pearl, J, 1990. "System Z: A natural ordering of defaults with tractable applications to default reasoning" *Proc. of the 2nd Conf. on Theoretical Aspects of Reasoning about Knowledge* (*TARK'90*) 121–135, San Francisco, CA, Morgan Karfman.

Shoham, Y, 1988. *Reasoning about Change* MIT Press.

Smets, P, 1988. "Belief functions" In: *Non-Standard Logics for Automated Reasoning* (P Smets, A Mamdani, D Dubois and H Prade, editors), 253–286, Academic Press.

Smets, P, 1990a. "The combination of evidence in the transferable belief model" *IEEE Trans. on Pattern Anal. Mach. Intell.* **12** 447–458.

Smets, P, 1990b. "Constructing the pignistic probability function in a context of uncertainty" *Un certainty in Artificial Intelligence 5* (M Henrion et al., editors), 29–40, North-Holland.

Smets, P, 1995. "Quantifying beliefs by belief functions: An axiomatic justification" In: *Proc of the 13th Inter. Joint Conf. on Artificial Intelligence* (*IJACI'93*), 598–603, Chambéy, France, August 28–September 3.

Smets, P and Kennes, R, 1994. "The transferable belief model" *Artificial Intelligence* **66** 191–234.

Artificial intelligence—a modern approach by Stuart Russell and Peter Norvig, Prentice Hall. Series in Artificial Intelligence, Englewood Cliffs, NJ.

Yet another introduction to artificial intelligence? Don't we have enough of these already? This is what I thought—before I held a copy of the book in my hands for the first time.

In fact, this book is different in many aspects from every other general AI book you may have seen before. First of all, it's unique in the broad coverage of topics. It (almost, see below) has it all: the book's 27 chapters cover problem solving and search, logic and inference, planning, probabilistic reasoning and decision making, learning, communication, perception and robotics. And in each section you will find an incredible amount of useful and totally up to date material that has never been included in other textbooks so far. There is a lot to learn, for beginners and for advanced AI people. You always wanted to know about Socratic reasoners, demotion, the upward solution property, coercion, policy iteration, PAC learning, adaptive belief networks, convolution, bigram models, the Viterbi algorithm, skeletonization, the horizon problem and the like? Then this is the right book for you.

Also the more standard parts give a lot of "nonstandard" material. The logic sections, for instance, not only give the typical introduction to propositional and first order logic together with the usual inference procedures, they also give many useful hints how to use first order logic to actually represent aspects of the real world including measures, time, actions, mental objects and