

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 probabilité" *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 Kaufman.
- 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, Chambéry, 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

the like, and they contain a lot of information about how to implement efficient logical reasoners. The section on uncertain knowledge contains an excellent introduction to probabilistic reasoning and belief networks. Moreover, it introduces decision theory covering topics like multiattribute utility functions, decision networks, sequential decision problems and dynamic decision networks. The section on learning, one of my favourites, present all sorts of approaches ranging from “subsymbolic” back-propagation learning in neural nets, genetic algorithms, decision tree learning, explanation based on learning to inductive logic programming, and it puts all these approaches into perspective.

The book also contains very valuable information about how specific approaches and techniques were used in real applications, how successful they were—or why they failed. The planning section is a particularly good example of this. With this information the reader gets a pretty good feeling of what can be done at the moment, and where the big problems are.

A further important aspect distinguishing this book from others is the common unifying perspective under which all the different approaches are presented. The authors view AI as the science of intelligent agent design. Under this view all the bits and pieces from various subfields of AI fit together very nicely. For novices this provides a lot of orientation. Advanced researchers get the great feeling that what they do is not only relevant to them and their little AI subcommunity.

The book is very well-written and clear with an excellent balance between motivation, formalisation and application. To make the underlying ideas precise the authors use easily understandable pseudo code throughout the book. Actual Common Lisp implementations of the presented algorithms are available via the Internet. The authors show a great ability to invent illustrating and entertaining examples—often reappearing in several chapters—and their style of writing is very amusing. It’s just fun to read this book.

After all this enthusiasm for the book, is there any wish left open? There is. I don’t want to mention the few mistakes one finds, rather unsurprisingly given this is the first edition of a book of over 900 pages (of course, consistency of first order logic is *not* semi-decidable, contrary to what is stated on p277). Number 1 on my wish list is a more adequate treatment of nonmonotonic reasoning. Much excellent work has been done in this area in recent years and interesting insights have been gained. All Russell and Norvig, basically, have to say about this is: theoretically interesting but practically irrelevant. I think this is too much of an oversimplification and the topic of nonnumerical defeasible reasoning deserves more than one page in a book like this. I would hope to see one or two extra chapters on this topic in a future edition.

Anyway, there can be no doubt that this is the best general AI textbook available today. If the quality of textbooks mirrors the maturity of a field AI is in much better shape than many of us may have thought. The book is a highly valuable source of information, not just for newcomers. Given its reasonable (if not cheap) price there is a pretty good chance that this will become the AI bible of the next decade.

Reviewed by Gerd Brewka, Technische Universität Wien, Austria

Concept formation and knowledge revision by Stefan Wrobel, Kluwer Academic, Netherlands, 1994, pp 240, £57.75, ISBN 079239500X.

The construction and maintenance of a knowledge base is a vast and sometimes confusing area of AI and machine learning, engulfed in specific theories and representations from various related fields.

Wrobel presents a holistic approach to the symbolic representation of knowledge, concept formation and updating methods for knowledge bases. The book has an interdisciplinary view and draws from close fields such as psychological studies on human learning, logical formalisms, computational learning and existing machine learning theories. In particular, the author’s system, MOBAL, bases its knowledge acquisition system on the architecture described in this book.