Book Review

A Review of *Nonmonotonic Reasoning*

Simon Parsons

It is possible to argue, relatively convincingly, that any research topic only begins to become mature when it appears on a syllabus somewhere. Once the topic has become well enough understood that it can be explained easily to paying customers, and stable enough that anyone teaching it is not likely to have to update his/her teaching materials every few months as new developments are reported, it can be considered to have arrived. Another reasonable indicator of the maturity of a subject, a milestone along the road to academic respectability, is the publication of a really good book on the subject—not another research monograph but a book that consolidates what is already known, surveys and relates existing ideas, and maybe even unifies some of them. Grigoris Antoniou's *Nonmonotonic Reasoning* is just such a milestone—well written, informative, and a good source of information on an important and complex subject.

Since the idea was first mooted around 1980,[1] there has been a vast amount of work under the banner of *nonmonotonic reasoning*, and it is neither surprising nor unreasonable that Antoniou only covers part of the territory. Neither is it surprising nor unreasonable that he devotes a lot of space to Reiter's (1980) default logic, which, along with McCarthy's (1980) circumscription and Moore's (1985) autoepistemic logic, is one of the holy trinity of nonmonotonic reasoning. Default logic is certainly the approach that most people find easiest to understand, and it has been the basis of a number of different variants, all with their own strengths and weaknesses. Antoniou not only discusses these variants in depth but also presents a new approach to computing the extensions of default theories that provides a means of relating the different approaches. The part of the book on default logic one might consider reasonably standard, and it is no surprise that it is followed by chapters on autoepistemic logic, circumscription (even if advocates of the theory might find the 14 pages on circumscription almost dismissively little), and the study of nonmonotonic consequence relations. However, at this point, *Nonmonotonic Reasoning* starts throwing in some more unusual topics.

Nonmonotonic Reasoning,
Grigoris Antoniou,
contributions by
Mary-Anne Williams,
The MIT Press,
Cambridge, Massachusetts, 1997,
285 pp.,
ISBN 0-262-01157-3.

The first of these more unusual topics is belief revision, which is covered in three chapters written by Mary-Anne Williams. I say *unusual* because approaches to belief revision are often considered to have rather different aims from those of most formalisms for nonmonotonic reasoning. Formalisms for nonmonotonic reasoning are, essentially, mechanisms for drawing reasonable conclusions from incomplete information, and formalisms for belief revision are mechanisms for dealing with overcomplete information, which includes contradictions. However, as Antoniou and Williams argue, techniques for belief revision can be considered dynamic approaches to nonmonotonic reasoning (because their reasoning certainly is nonmonotonic) in contrast to the static approaches of systems such as default logic. Even this distinction, though, is rather artificial because even default logic and its brethren obtain their nonmonotonicity from the fact that they can handle inconsistency by retracting old assumptions (even if this retraction is, in practice, achieved by recomputing conclusions), and *Nonmonotonic Reasoning* is considerably richer for the inclusion of the material on belief revision.

The second topic that strikes me as somewhat unusual is that of logic programming. Again, I find this unusual more because it is conventional to think of the study of nonmonotonic logics and the study of logic programming as different, being served by different conferences and having different agendas, than because there is something fundamentally wrong in considering the two areas as complementary. In addition, of course, the two areas complement each other rather nicely, with work on nonmonotonic logics stressing representation issues, what it is possible to capture in nonmonotonic reasoning, and logic programming stressing sound computation, what it is possible to compute and what the semantics of this computation are. As Antoniou points out, these concerns overlap, for example, in the stable semantics for logic programs, which is basically a default logic interpretation. Because of this slight difference in approach between logic programming and nonmonotonic logic, I think that this section of the book is a valuable addition. I do feel, however, that the book would have been stronger if it had covered the recent

argumentation semantics for logic programs. In my view, the argumentation approach is important because it can also be used to underpin nonmonotonic logics (Bondarenko et al. 1997), thus providing a unified approach to most of the material in *Nonmonotonic Reasoning*. However, this omission is fairly minor.

In fact, all the things, and there aren't many, that I don't like about *Nonmonotonic Reasoning* are rather minor. The most numerous of these are quibbles about its style (a fact that in itself is a recommendation). First, although each chapter gives a number of relevant problems for the reader, no solutions are given. Now, it might be argued that these are problems for which there is no particular solution (many are of the form "complete the proof of this theorem from the text"); so, giving such a solution is pointless. However, I disagree. If you set problems and give no solutions, you are abandoning those students interested enough to want to attempt the problems but incapable of solving them unaided (as someone who used to be exactly this kind of student, I get annoyed by the memory of all the times it happened to me). Second, a number of the proofs in the book are left incomplete (and the completion is often left as an exercise for the reader). I find this annoying for exactly the same reason as I find exercises without answers annoying—What are you meant to do when you can't complete the proof yourself? Finally, and as a criticism this is the weakest of the lot, I find the text rather dry in places, with a lot of mathematics unleavened with much explanation of the significance of the results. I always find it rather disappointing when research monographs suffer from this problem (surely when writing a book, one has the time and space to expand a little on the significance of the mathematics, to sing and dance a little in celebration of what it all means?), and in an introductory text, I find it doubly disappointing.

However, despite these rather minor problems, I think that *Nonmonotonic Reasoning* is a very fine book and one that I have no hesitation in recommending to anyone who wants to get a grip on this important area.

Note

1. Not having become aware of the concept until much later, I can't put an accurate date on the birth of nonmonotonic reasoning. A convenient landmark is the publication of volume 13 of *Artificial Intelligence* in 1980 because it contained a number of landmark papers on the subject, but some of the ideas were clearly knocking around for several years beforehand.

References

Bondarenko, A.; Dung, P. M.; Kowalski, R. A.; and Toni, F. 1997. An Abstract Argumentation-Theoretic Approach to Default Reasoning. *Artificial Intelligence* 93(1): 63–101.

McCarthy, J. 1980. Circumscription—A Form of Nonmonotonic Reasoning. *Artificial Intelligence* 13(1): 27–39.

Moore, R. 1985. Semantical Considerations on Nonmonotonic Logic. *Artificial Intelligence* 13(1): 75–94.

Reiter, R. 1980. A Logic for Default Reasoning. *Artificial Intelligence* 13(1): 81–132.

Simon Parsons is currently a lecturer in the Intelligent Systems Applications Group, part of the Department of Electronic Engineering at Queen Mary and Westfield College, London. He has a B.A. in engineering from the University of Cambridge; an M.Sc. in computing science from the Imperial College of Science, Technology, and Medicine; and a Ph.D. in AI from Queen Mary and Westfield College. His research interests center on techniques for decision making with imperfect information in intelligent systems; his current focus is on the applications of such techniques to risk assessment. His e-mail address is S.D.Parsons@elec.qmw.ac.uk.