Book review

Explanatory Nonmonotonic Reasoning by Alexander Bochman World Scientific,
Hardback: ISBN 981-256-101-3, Price: $68.00, xiv + 408 pages.

1 Introduction

This book is written with immense erudition and an unusual, almost superhuman
knowledge of the area. Unlike other parts of mathematical logic, and for a variety of
reasons, the number of researchers of the area of nonmonotonic logic was (and may
still be) immense. Since they all have to publish (or perish), there is a very large body
of the literature of the topic. It appears the author of this book, Alexander Bochman,
read and (even more importantly) digested those papers. In this process he created
a perspective of the area. This is the topic of the reviewed book. So, in my view, the
author is, probably, the most knowledgeable individual in the area of nonmonotonic
reasoning. The reviewed book contains an immense amount of material covering all
the fundamental modes of reasoning in nonmonotonic reasoning (NMR): default
logic; logic programs with a variety of semantics (primarily stable semantics, but
also Clark's completion, and multivalued semantics), argumentation theory, modal
nonmonotonic logics, and also circumscription (minimal models). The breadth of
the material in the book is huge. I would venture that most (if not all) significant
results of NMR as we know it in 2006 can be found in this book. Those of us who
tried to give a unified perspective of the area[1] know how difficult is it to put all the
information (or the most of what is significant in it) found in the past 25 years in a
unified coherent picture.

First, what it is all about? Starting in the 1960s, with the pioneering work by
J. McCarthy and then in the late 1970s with the work of late R. Reiter on CWA
continued by J. Doyle, J. Minker and D. McDermott it became clear that the long
(length measured in thousands of years) held belief, that any valid and useful method
of reasoning must be (essentially) classical logic, was incorrect. Of course, once the
dust started to settle, it became clear that philosophers and mathematical logicians
had dealt with nonmonotonicity in some special cases (for instance by limiting the
semantics of formulas). Yet the idea that some fragments of commonsense reasoning
(for instance the "concluding by default") can, sometimes at least, be formalized was
quite revolutionary. We will spare the reader some juicy tidbits of the developments
in the area. It will be sufficient to say that the topic was originally quite controversial.
The NMR as an area of research erupted with the special issue on NMR in the
Artificial Intelligence Journal in 1980 with important results by Reiter, McCarthy,

---

[1] Full disclosure: the author of this review tried this (together with M. Truszczyński) almost 15 years
ago.

Doyle and McDermott. By 1985, after the introduction of Autoepistemic Logic by R. Moore, the area was reasonably well delineated. In parallel there were efforts in a different community (logic programming) to understand the semantics of programs admitting negation in the body. Actually, these efforts started earlier. Yet another community that contributed to the area was the theoretical database community. All these loosely related (but non-disjoint) communities influenced each other.

It should be observed that nonmonotonicity occurs in a large number of areas of Computer Science, Mathematics and Philosophy. One relatively close area is studying nonmonotonicity of axiomatically-defined provability relations. This area is only tangentially discussed in the reviewed book.

The explanatory nonmonotonic reasoning of the title is based on the following schema (as we will see with a number of degrees of freedom). I am explaining things as I see them, the author may be looking from some other perspective. In this schema there are some symbolic objects (formulas, maybe sets of formulas, maybe even collections of sets of formulas) and a context S. This context may be one set of formulas or several such sets. There is some abstract provability relation. There is a way of filtering with the context S of the syntactic objects under consideration. After the filtration there is a fixpoint computation. If the result of the computation coincides with the fixpoint, this context S is an *extension* (additional tests may be involved, for instance some minimality checking). The presence of the formula in the fixpoint is 'explained' by some proof-theoretic process. For instance, in the case of stable models the explanation involves collecting negative information (which is collectively captured by Gelfond-Lifschitz reduction) which explains why a given atom is in the least model of $GL(P, M)$ in terms of atoms that must be *absent* from M. If this sounds like Reiter's extensions, or Gelfond-Lifschitz stable models, or answer sets, of Fitting's four-valued stable models, or some other familiar construction, it is not an accident. This is what the technology developed in this book tries to capture[2]. As we see there are at least four, if not five "parameters" that can be plugged in: syntactic objects, underlying logical language, filtration process, closure, and possibly additional check. In such abstract setting we can expect that many modes of reasoning will be represented, and indeed they are. In Bochman's proposal the basic underlying structure is called a *bisequent*. A bisequent $a : b \Vdash c : d$ (where a, b, c, d are finite sets of objects (but the intuition is that the objects are propositions of some formal language) has the following interpretation (in fact there are several, we will quote two of p. 58 of the reviewed book). '*If no proposition from b is assumed, and all propositions from d are assumed, then all propositions in a hold only if at least one of propositions in c holds*'. Yet another, and closer to the intuitions held by this reviewer, is this one: '*If all propositions from a are assumed and no proposition from b is assumed then either (at least one) proposition from c holds, or one of the propositions of d is false*". This last reading is in the spirit of Lifschitz and Woo, namely as a general program rule. In this context, let us mention that the answer sets of the kind proposed by Gelfond and

---

[2] Many years ago Truszczyński and the author tried to capture a simplified version of this scheme, hence the way I see this proposal.

Lifschitz in their work on Logic Programming (and thus also Lifschitz and Woo), in particular disjunctive logic programming involve the additional test mentioned among 5 degrees of freedom in the scheme discussed above.

It is also clear (and it has been discovered by Fitting and Kunen in the 80ies) that 3- and 4-valued logic (but of Kleene, and of Belnap variety, not of Post family of logics) are particularly well-suited for such considerations.

2 Contents

Now, to the actual contents of the book. The first hurdle that needs to be overcome is the terminology. In his quest to put, more or less, the entire area into the uniform approach, the author created his own language (the bisequents). But most of the community (maybe the entire community) did not adopt this language. This creates a serious problem for the reader.

The author introduces the nonmonotonic reasoning in Chapter 1. There is (I guess, to gain a favor with philosophers?) an obligatory citation from Wittgenstein (brr …) which, in my view sets up the tone of the book. On a serious side, the author introduces the motivation and outlines the scope of the book. The two strands of NMR (explanatory, more or less coinciding with the description above), and one based on studies of abstract consequence relations are distinguished. To some extent the relation of NMR and the underlying computational mechanism is explored here (but not enough to my taste).

Chapter 2 introduces the reader to several important issues. Scott- and Tarski-consequence relations are explored. The issues of minimality and of support are discussed.

This naturally leads to the introduction of the completion and of loop-formulas (that capture the difference between the stable and supported semantics). Some elements of circumscription (a major topic are in itself) are studied.

Four-valued logic of Belnap and its interpretations are discussed in the next chapter. Here, the four-valued logic is treated as a composition of two-copies of two-valued logic (the idea due to Ginsberg and Fitting). There is plenty of detail and a significant body of knowledge presented here, and in quite detail.

There is a natural interpretation of default logic and of logic programs as bisequents. This is the topic of the next chapter, Nonmonotonic Semantics. Several important results, including those on strong equivalence of programs are presented there. Let me mention that (with the work of Pierce, Lifschitz and collaborators on the connections to the logic of "here and there" and of Eiter, Turner and Truszczyński, both on variations of that concept, and on alternative presentations) we now have a very clear understanding of the phenomena of strong equivalence. Finally, one sees that bisequents are, in reality, program clauses in disguise (if one treats the formulas as some kind of atoms tied with other atoms via semantics). The connections with the stationary semantics of Przymusinski are shown.

A thorough investigation of default logic (extensions and other structures associated with default theories) is the subject of the next chapter. It is a well-known fact that, essentially, stable semantics of Gelfond and Lifschitz is very naturally

interpretable in default logic. In fact, modulo proper modeling of rules of logic, there is no difference whatsoever. But there are (besides an obvious interpretation as a very simple bisequent) several others, and they are neatly classified in this chapter.

The bisequent data structure allows for a natural mode of nonmonotonic theory of argumentation (Bondarenko, Dung and Kowalski). The next chapter is devoted to this theory.

Both default logic and stable semantics have been discovered in the quest for the solution to the frame problem. In fact, the major explanatory success of stable semantics of logic program is its solution to the frame problem. Indeed, the kind of negation stable semantics provides is – as shown by Gelfond and Lifschitz – precisely the one that is needed for successful formalization of the frame problem. This leads to the studies of causal relationship and the use of logic programming for providing a correct formalization of causality which is the subject of the next chapter. This area still awaits a definitive presentation and is the subject of ongoing research.

The two final chapters are devoted to epistemic (and thus Kripke-style) semantics for NMR, and to modal nonmonotonic logics. For a variety of reasons (similarly to the argumentation theory) this area is out of fashion lately. But the truth is that modal formalizations of NMR allow for capturing many (if not all) aspects of nonmonotonic logics. There are several interesting special cases corresponding to Segerberg's maximal logics (below S5): those are Moore's Autoepistemic Logic, and Schwarz' Reflexive Autoepistemic Logic. They are discussed in detail. A natural nonmonotonic modal logic corresponding to default logic, Truszczynski's nonmonotonic S4F (with its relation to Gödel's interpretation of intuitionistic logic in modal logic), is also presented and studied.

3 Conclusions

The short description of the contents of the reviewed book implies, we hope, that the reader now see why we put forward our original assessment of the unusual depth and richness of this book. We believe that most technical results in the literature and many results that will be published by many researchers, are in this book already. So now, the question arises if this book could be used as a handbook in some graduate course of Knowledge Representation and Reasoning. To this question my answer would be an unequivocal "no". If the reader does not already have a working knowledge of NMR, it is more than likely she will not learn it from there. There are several reasons for this (quite sad) conclusion. First, there are very few examples. I would say that most abstract papers in the obscure areas of, say, Category Theory, are written in this way. But NMR has roots in knowledge representation and reasoning. It was invented to solve the practical problems rooted in practical formalized reasoning. These roots almost completely disappeared from the picture the author delivers. Next, from early days in the history of NMR (but certainly from the beginning of the 1990s – viz. papers by Nerode, Remmel and this reviewer, and even more urgently in several papers by Niemelä, and Truszczynski and this reviewer) it became clear that NMR is also about computing. These aspects

of NMR, today a vibrant area called *Answer Set Programming* are simply absent from the reviewed book[3]. Of course the author had to choose the material, and it is an indisputable right of the author to push his own message. But it should be clear from this review that I do not agree with what he chose. Since Herbrand, Hilbert and Gödel we know that logic is, at least to some extent, a form of computation. In logic, as we compute the consequences, we *compute* a fixpoint. It is plain symbolic for the point of view presented in the reviewed book, that there is no place for Tarski fixpoint theorem in the bibliography of the reviewed text.

Time for the summing up. I will not suggest to any novice, and any graduate student of NMR, that they use this book to learn the fundamentals of NMR. They will not learn it from there. But I will gladly keep this book on my shelf and, I am sure, I will often look it up if some argument, in broadly understood NMR, will be needed. It is likely I will find it there.

Victor W. Marek
Department of Computer Science
University of Kentucky

---

[3] From this perspective, the cover of the book, referring to computation is especially misleading.