

## Book Review

RONALD FAGIN, JOSEPH Y. HALPERN, YORAM MOSES AND MOSHE Y. VARDI  
*Reasoning About Knowledge*. MIT Press. 1995. ISBN 0-262-06162-7. \$45.00. xiii+477pp.  
Hardbound.

Reasoning about knowledge is an exciting field of contemporary logic, placed as it is at the crossroads of such diverse disciplines as artificial intelligence (AI), economics, cryptography, pragmatics of natural language, database theory and distributed systems theory. It should not be confounded with that subfield of AI which goes under the name of “knowledge representation and reasoning”. Although there is a substantial overlap of concerns between the two, and although both deal with the logical analysis of agents’ knowledge, only the former presupposes an *explicit* representation of the notion of “knowledge”, usually accomplished by means of operators fashioned on those of modal logic. Reasoning about knowledge is therefore that discipline concerned with analysing the implications of arguments such as “Bob does not know whether he has mud on his forehead, but he knows that Alice knows it, and that she knows he does not know”.

It cannot be claimed that this is a comprehensive handbook of reasoning about knowledge, as it is mainly devoted to the “semantic” approach (a rival account of these notions, initiated by Montague and not covered in the present book, is based on the “syntactic” treatment of epistemic modalities). However, this is by no means a substantial limitation, as the possible worlds approach is the one that has been most fruitful and proved most flexible in analysing knowledge and a number of applicatively interesting, related notions. The 11 chapters of the book embrace the “epistemic” concepts of knowledge and belief of multiple agents (either natural or artificial), their relationships with time and action, their model- and proof-theoretic formalisations, their decidability and computational complexity properties, their numerous applicative contexts. This work is presented within the unifying framework of Hintikka-Kripke’s possible worlds semantics, presented in Chapter 2 in its simplest form and progressively extended throughout the book as more and more complex notions are analysed.

No-one could be better qualified than these authors to write a textbook on this topic. Key contributors to the literature on reasoning about knowledge and its applications within AI, distributed systems and cryptography, they are to be credited for much of the recent development of this discipline and its growth in importance outside the mere philosophical logic circles. The book is written in a clear, terse style; the presentation and the mathematical arguments are both rigorous and detailed, and the close interplay of formal mathematical work and informal philosophical argument makes the book very pleasant to read. The book is as self-contained as a logic textbook could be hoped to be, requiring only basic knowledge of logic methodology. In recent years, I have used some of this material (some had already been published in journal articles) while teaching undergraduate courses on logic and AI, and have found it of great pedagogical value, not least for its exemplary way of proceeding through the steps of a logical argument and formalisation. Each chapter is completed with a set of exercises and a comprehensive annotated bibliography.

It is easy to foresee that this book will become a classic.

FABRIZIO SEBASTIANI

*Istituto di Elaborazione dell'Informazione  
Consiglio Nazionale delle Ricerche  
56126 Pisa, Italy*