The idea that defines the very heart of "traditional" Artificial Intelligence (AI) is due to John McCarthy: his imagined ADVICE-TAKER was a system that would decide how to act (in part) by running formal reasoning procedures over a body of explicitly represented knowledge, a *knowledge base*. The system would not so much be programmed for specific tasks as told what it needed to know, and expected to infer the rest somehow. Knowledge and advice would be given declaratively, allowing the system to operate in an undirected manner, choosing what pieces of knowledge to apply when they appeared situationally appropriate. This vision contrasts sharply with that of the traditional programmed computer system, where what information is needed and when is anticipated in advance, and embedded directly into the control structure of the program.

This is a book about the *logic* of such knowledge bases, in two distinct but related senses. On the one hand, a knowledge base is a collection of sentences in a representation language that entails a certain picture of the world represented. On the other hand, *having* a knowledge base entails being in a certain state of knowledge where a number of other epistemic properties hold. One of the principal aims of this book is to develop a detailed account of the relationship between symbolic representations of knowledge and abstract states of knowledge.

This book is intended for graduate students and researchers in AI, database management, logic, or philosophy interested in exploring in depth the foundations of knowledge, knowledge bases, knowledge-based systems, and knowledge representation and reasoning. The exploration here is a mathematical one, and we assume some familiarity with first-order predicate logic (and for motivation at least, some experience in AI).

The book presents a new mathematical model of knowledge that is not only quite general and expressive (including but going well beyond full first-order logic), but that is much more workable in practice than other models that have been proposed in the past. A reader can expect to learn from this book a style of semantic argument and formal analysis that would have been quite cumbersome, or even outside the practical reach of other approaches.

From a computer science point of view, the book also develops a new way of specifying what a knowledge representation system is supposed to do in a way that does not make assumptions about how it should do it. The reader will learn how to treat a knowledge base like an *abstract data type*, completely specified in an abstract way by the knowledge-level operations defined over it.

The book is divided into two sections: Part I, consisting of Chapters 1 to 8, covers the basics; Part II, consisting of Chapters 9 to 14, considers a number of more-or-less independent research topics and directions. (The contents of these chapters are described at the end of Chapter 1.) The material in the book has been used in graduate level courses at the authors' institutions in Canada and Germany. In one semester, it should be possible to cover all of Part I and at least some of the advanced chapters of Part II. Exercises and bibliographic notes are included at the end of each chapter. Suggestions for further research are made at the end of the chapters of Part II. An index of the important technical terms, whose first use is underlined in the text, appears at the end of the book. Comments and corrections are most welcome and can be sent to the authors at

hector@cs.toronto.edu gerhard@cs.rwth-aachen.de.

Although every effort has been made to keep the number of errors small, this book is offered as is, with no warranty expressed or implied.