

Preface

Computational Intelligence: A Logical Approach is a book about a new science—computational intelligence. More commonly referred to as artificial intelligence, computational intelligence is the study of the design of intelligent agents. The book is structured as a textbook, but it is accessible to a wide audience.

We wrote this book because we are excited about the emergence of computational intelligence as an integrated science. As with any science worth its salt, computational intelligence has a coherent, formal theory and a rambunctious experimental wing. We balance theory and experiment, and show how to link them intimately together. We develop the science of CI together with its engineering applications. We believe the adage, “There is nothing so practical as a good theory.” The spirit of our approach is captured by the dictum, “Everything should be made as simple as possible, but not simpler.”

The book works as an introductory text on artificial intelligence for advanced undergraduate or graduate students in computer science, or related disciplines such as computer engineering, philosophy, cognitive science, or psychology. It will appeal more to the technically-minded; parts are technically challenging, focusing on learning by doing: designing, building, and implementing systems. Any curious scientifically-oriented reader will profit from studying the book. Previous experience with computational systems is desirable, but prior study of logic is not necessary, as we develop the concepts as required.

The serious student will gain valuable skills at several levels ranging from expertise in the specification and design of intelligent agents to skills for implementing, testing, and improving real software systems for several challenging application domains. The thrill of participating in the emergence of a new science of intelligent agents is one of the attractions of this approach. The practical skills of dealing with a world of ubiquitous, intelligent, embedded agents are now in great demand in the marketplace.

The focus is on an intelligent agent acting in an environment. We start with simple agents acting in simple, static environments and gradually increase the power of the agents to cope with more challenging worlds. We make this concrete by repeatedly illustrating the ideas with three different agent tasks: a delivery robot, a diagnostic assistant, and an information slave (the infobot).

Our theory is based on logic. Logic has been developed over the centuries as a for-

mal (that is, precise not obtuse) way of representing assumptions about a world and the process of deriving the consequences of those assumptions. For simple agents in simple worlds we start with a highly restricted simple logic. Then as our agent/environment requires, we increase the logical power of the formalism. Since a computer is simply a symbol-manipulation engine, we can easily map our formal theories into computer programs that can control an agent or be used to reason about an agent. Everything we describe is implemented that way.

Any high-level symbol-manipulation language, such as LISP or Prolog, can be used to support these implementations. We also provide complete implementations in Prolog, a language that is designed around the idea of computations as proof in a simple logic. But this connection is not essential to an understanding or use of the ideas in this book.

Our approach, through the development of the power of the agent's representation language, is both simpler and more powerful than the usual approach of surveying and cataloging various applications of artificial intelligence. However, as a consequence, some applications such as the details of computational vision or the pragmatics of low-level robot control are not covered in this book. We do show how our theory allows us to design situated robots performing in dynamically changing worlds.

We invite you to join us in an intellectual adventure: building a science of intelligent agents.

David Poole
Alan Mackworth
Randy Goebel

Acknowledgments

We are very grateful to Valerie McRae for contributing her proofreading and formatting skills and for helping in many other ways. We also gratefully acknowledge the helpful comments on earlier drafts of this book received from Carl Alphonse, Brock Atherton, Veronica Becher, Wolfgang Bibel, Craig Boutilier, Jason Burrell, Hugh Clapin, Keith Clark, Jim Delgrande, Koichi Furukawa, Chris Geib, Holger Hoos, Michael Horsch, Ying Li, Greg McClement, Claude Sammut, Julie Schlifské, Daniele Theseider Dupré, Paul van Arragon, Bonnie Webber, Bob Woodham, and Ying Zhang. Thanks to our editors at OUP, Bill Zobrist, Karen Shapiro, and Krysia Bebick.

The four-limbed robot shown on the cover is a *Platonic Beast*, developed by Dinesh Pai, Rod Barman, and Scott Ralph at the University of British Columbia's Laboratory for Computational Intelligence. Based on the spherical symmetry of the tetrahedron, one of the five symmetrical Platonic polyhedra, the Beast can walk, roll over, and climb over obstacles. Aidan Rowe and John Chong provided the technical graphical expertise to complete the cover design.