PREFACE

The goal of the M.I.T. Press Series in Artificial Intelligence is to make available to a wide and interdisciplinary audience some of the work that is occuring at the frontiers of artificial intelligence research. In many cases, this means that the books in the series will be snapshots of work in progress, rather than the rounded, well-polished presentations that characterize older, more stable fields.

This book is such a snapshot. It is a slightly edited version of my Ph.D. thesis, completed in September of 1977 at the M.I.T. Artificial Intelligence Laboratory. It assumes that the reader is familiar in general with the principles of digital computers and has some appreciation for the difficulties involved in making these devices behave intelligently. It does not assume a knowledge of particular details or languages, except in the section on software simulation of the knowledge network and the appendices on hardware implementations. It is my hope that this book will be of interest to psychologists, linguists, and others not directly connected to the fields of computer science or artificial intelligence.

I would like to express my gratitude to the following people, who made essential contributions to this project:

To my thesis supervisor, Gerald Sussman, and my readers, Patrick Winston and Marvin Minsky, for their advice, support, and patience. I would also like to thank each for the special gifts that he brought to the task: Sussman, his contagious energy and enthusiasm; Winston, his broad perspective and good sense; and Minsky, his sense of direction.

To Drew McDermott, Bob Moore, Michael Genesereth, Mitch Marcus, and Jon Doyle for the many stimulating arguments and discussions that we have had about knowledge-representation and related topics.

To McDermott, Genesereth, Doyle, Candy Bullwinkle, Bill Martin, Guy Steele, Steve Rosenberg, Jim Stansfield, and

Johan de Kleer for reading drafts of the thesis and for making many useful comments and criticisms.

To Tom Knight, Jack Holloway, Kurt VanLehn, Ben Kuipers, David Marr, Carl Hewitt, Bob Woodham, Richard Stallman, Richard Greenblatt, Brian Smith, Chuck Rich, and Howie Shrobe for contributing good ideas, for pointing out things that I ought to read, for providing tools that were a pleasure to work with, for general moral support, and for combinations of the above.

To Bob Sjoberg and Brian Reid for their help in the area of digital calligraphy.

Finally, and most of all, I would like to thank my wife, Penny, for her support and understanding. This book is dedicated to her.

For two of the four years covered by this research, I was supported by an I.B.M. Graduate Fellowship. The laboratory was funded in large part by the Defense Advanced Research Projects Agency, as noted below.

Scott Fahlman

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the Laboratory's artificial intelligence research is provided in partby the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643.