
Preface

Artificial intelligence, or AI, is largely an experimental science—at least as much progress has been made by building and analyzing programs as by examining theoretical questions. MYCIN is one of several well-known programs that embody some intelligence and provide data on the extent to which intelligent behavior can be programmed. As with other AI programs, its development was slow and not always in a forward direction. But we feel we learned some useful lessons in the course of nearly a decade of work on MYCIN and related programs.

In this book we share the results of many experiments performed in that time, and we try to paint a coherent picture of the work. The book is intended to be a critical analysis of several pieces of related research, performed by a large number of scientists. We believe that the whole field of AI will benefit from such attempts to take a detailed retrospective look at experiments, for in this way the scientific foundations of the field will gradually be defined. It is for all these reasons that we have prepared this analysis of the MYCIN experiments.

The MYCIN project is one of the clearest representatives of the experimental side of AI. It was begun in the spring of 1972 with a set of discussions among medical school and computer science researchers interested in applying more intelligence to computer programs that interpret medical data. Shortliffe's Ph.D. dissertation in 1974 discussed the problem and the MYCIN program that implemented a solution. In itself, the 1974 version of MYCIN represents an experiment. We were testing the hypothesis, advanced in previous work at Stanford, that a rule-based formalism was sufficient for the high performance, flexibility, and understandability that we demanded in an expert consultation system. The positive answer to this question is one of the best-known lessons in the history of AI.

In addition to, or rather because of, the original MYCIN program and the medical knowledge base that was accumulated for that work, many derivative projects explored variations on the original design. EMYCIN¹ is among the best known of these, but there are several others. In this book we discuss many of the experiments that evolved in the period from 1972

¹We use the name EMYCIN for the system that evolved from MYCIN as a framework for building and running new expert systems. The name stands for "essential MYCIN," that is, MYCIN's framework without its medical knowledge base. We have been reminded that E-MYCIN is the name of a drug that Upjohn Corp. has trademarked. The two names should not be confused: EMYCIN should not be ingested, nor should E-MYCIN be loaded into a computer.

to 1982 based on the 1972–1974 design effort. We have chosen those pieces of work that, at least in retrospect, can be seen as posing clear questions and producing clear results, most of which were documented in the AI or medical literature and in technical reports.

We are taking a retrospective view, so as to restate questions and reinterpret results in a more meaningful way than that in which they were originally documented. Among other things, we now present these pieces of work as a collected whole, whereas they were not originally written as such. Each paper is heavily edited—new sections have been added to put the work in context, old sections have been deleted to avoid redundancies and “red herrings,” and the entire text has been reworked to fit each paper into the unified picture. Each part begins with an overview chapter posing the central question of the section, discussing the implications of the question in its historical context, and providing a current framework for interpreting the results. Some entirely new papers were prepared specifically for this book. In addition, we are including several papers and technical reports that have previously been difficult to find and will therefore be generally available for the first time.

The last chapter is entirely new and could not have been written until the experiments were performed. It presents a set of conclusions that we have drawn from the experimental results. In a sense, the rest of the book discusses the data that support these conclusions. We believe this book is unique in its attempt to synthesize 10 years of work in order to demonstrate scientific foundations and the way in which AI research evolves as key issues emerge.

Acknowledgments

We gratefully acknowledge the help and friendship of Edward Feigenbaum and Joshua Lederberg. Over many years not only did they motivate and encourage our work in applying artificial intelligence to medicine, but they also created the intellectual and computing environment at Stanford that made this work possible.

Many individuals have contributed to the varied aspects of MYCIN and to the ideas in this book. In the past we have affectionately referred to each other as “the MYCIN gang.” All of the authors of chapters built parts of MYCIN, performed experiments, contributed to overall design, and/or wrote the original articles on which the chapters are based. The persons who have been part of the MYCIN gang with us over the years are Janice S. Aikins, Stanton G. Axline, Timothy F. Beckett, James S. Bennett, Sharon Wraith Bennett, Robert L. Blum, Miriam B. Bischoff, Alain Bonnet, A. Bruce Campbell, Robert Carlson, Ricardo Chavez-Pardo, William J. Clancey, Jan E. Clayton, Stanley N. Cohen, Gregory F. Cooper,

Randall Davis, Robert S. Engelmores, Lawrence M. Fagan, Robert Fallat, Edward A. Feigenbaum, John Foy, Jean Gordon, Cordell Green, John F. Hannigan, Diane Warner Hasling, Robert Illa, Charlotte D. Jacobs, Jonathan J. King, John C. Kunz, Reed Letsinger, Robert London, Dana Ludwig, Thomas C. Merigan, John J. Osborn, Frank Rhame, Louis Sanner, A. Carlisle Scott, David E. Smith, Motoi Suwa, Randy L. Teach, William J. van Melle, Jerold Wallis, and Victor L. Yu.

We wish to thank, also, the numerous other persons in the AI, particularly the AIM (AI in Medicine), community with whom conversations over the years have been valuable sources of ideas. In particular, Randy Davis, John McDermott, Carli Scott, Bill van Melle, Bill Clancey, and Jim Bennett gave us help in clearly formulating Chapter 36.

We were fortunate in having Joan Differding and Dikran Karaguezian organize and edit much of the material here. Their help was invaluable. Jane Hoover carefully copyedited the manuscript and substantially contributed to its readability. We also wish to thank Darlene Vian, Juanita Mullen, Susan Novak, Cindy Lawton, and Barbara Elspas for assistance in preparing the manuscript.

Almost all of the computing work reported here (and the manuscript preparation) was done on the SUMEX-AIM computer at Stanford. Mr. Thomas Rindfleisch, former director of SUMEX, made it possible for us to attend to the research without undue worry about system reliability by making SUMEX a stable, high-quality computing environment. We also wish to acknowledge the importance of SUMEX funding from the National Institutes of Health, Division of Research Resources (RR-00785) and to thank, particularly, Dr. William Baker for his leadership in sustaining the quality of the resource.

Finally, the individual projects described here were funded by a variety of agencies over the years. We gratefully acknowledge the assistance of the Bureau of Health Services Research and Evaluation (HS-01544), the National Institutes of Health, General Medical Sciences (GM 01922 and GM 29662), the National Science Foundation (MCS-7903753), the Defense Advanced Research Projects Agency (MDA-903-77-C-0322), the Office of Naval Research (N0014-79-C-0302, NR-049-479), the National Library of Medicine (LM-03395, LM-00048), and the Henry J. Kaiser Family Foundation.

Stanford University
March 1984

B.G.B.
E.H.S.