

Preface

This text is a description of a computer-based system designed to assist physicians with clinical decision-making. This system, termed MYCIN, utilizes computer techniques derived principally from the subfield of computer science known as artificial intelligence (AI). MYCIN's task is to assist with the decisions involved in the selection of appropriate therapy for patients with infections.

MYCIN contains considerable medical expertise and is also a novel application of computing technology. Thus, this text is addressed both to members of the medical community, who may have limited computer science backgrounds, and to computer scientists with limited knowledge of medical computing and clinical medicine. Some sections of the text may be of greater interest to one community than to the other. A guide to the text follows so that you may select those portions most pertinent to your particular interests and background.

Chapter 1, "Introduction" provides an introduction to the fields of medical computing, artificial intelligence, and the clinical problem area for which the MYCIN program is designed. It concludes with an introductory overview of MYCIN and a sample interactive session.

Chapter 2, "Design Considerations for MYCIN," presents the design criteria that were considered during MYCIN's development; acceptability by physicians is emphasized. The chapter ends with a brief discussion of how MYCIN attempts to satisfy the criteria.

Chapter 3, "Consultation System," describes in detail how the MYCIN program makes decisions. The data structures and control structures are discussed in the context of prior work regarding

rule-based problem-solving. Certain subsections of this chapter have been isolated and marked with an asterisk (see below) so that noncomputer scientists can read the more descriptive information without becoming overly immersed in the details of implementation.

Chapter 4, "Model of Inexact Reasoning in Medicine," covers the general topic of inferential model building and is a somewhat separate topic. It has therefore been written to be self-contained. If your primary interest is in MYCIN's truth model, you may concentrate on Chapter 4 without needing to refer to other parts of the book for clarification of details.

Chapter 5, "Explanation System," discusses MYCIN's ability to answer questions regarding both its knowledge base and the details of a specific consultation. The user interface is described in detail, but implementation information is omitted and may be found elsewhere [Shortliffe, 1974b].

Chapter 6, "Future Directions for MYCIN," introduces the several plans for evaluation and future extensions of the MYCIN system. These include immediate plans for working on knowledge acquisition procedures, and eventual implementation of the program as one module in a total Hospital Information System.

Finally, Chapter 7, "Conclusion," summarizes the program's accomplishments to date and discusses MYCIN's limitations, plus its contributions to the fields of computer-based medical decision making and artificial intelligence.

Each of the seven chapters is named with an arabic numeral and divided into sections, also specified by arabic numeral designations. Subsections are likewise named. Sections and subsections that are followed by an asterisk (*) denote those parts of the text that contain detailed technical explanations of MYCIN system components; these sections may be skipped. Sections followed by double asterisks (**) are summary sections and may be omitted if you are already well-acquainted with the topic. Reference citations are enclosed in square brackets (e.g., [author, 1976]) and cite the name of the first author plus the year of publication. If an author published more than one referenced article in a single year, a lower-case letter is appended to the date. An alphabetized bibliography appears at the end of the text.

A final point should be made regarding the use of the male pronoun to refer to physicians and patients throughout this work. I have decided to follow convention rather than inject awkwardness in

Preface

an effort to recognize both sexes. It therefore seems wise to stress from the outset that, although such a convention is less than ideal, 'he', 'him', and 'his' are meant to be interpreted without any gender association.

Acknowledgments: This text would still be in its infancy if it were not for the active participation and advice of several individuals. Particularly important to the work was the close association of a group of physicians and computer scientists who have met regularly for several years. At these collaborative sessions Stanley Cohen regularly provided invaluable contributions, and he also was the first person at Stanford to encourage me to extend my time in medical school in order to undertake this research. I am grateful for both his time and his help and guidance. Bruce Buchanan provided computer science expertise, a philosopher's outlook, and insightful observations without which this project might well have floundered. The analysis in Chapter 4 is largely a product of readings he suggested and mutual discussions of the philosophical issues involved. Stanton Axline also participated in the evolution of MYCIN. He provided the infectious disease expertise needed for the development of the program's knowledge-base and also suggested several design characteristics that have been implemented to heighten the system's acceptability to physicians. In late 1973 an additional researcher, Randy Davis, joined the project. Our many conversations have proved invaluable, particularly because he encouraged me to search for ways to generalize MYCIN's problem-solving approach.

Several other physicians have assisted during the initial development period. The project was conceived during idea sessions at which Thomas Merigan was an active participant, and Gilbert Hunn met regularly with us during the first year. I am also grateful to Michael Podlone and Robert Illa for their assistance testing the program on sample patient cases and for their suggestions regarding its performance. In addition I would like to thank Frank Rhame, Patrick Goodall, Richard Greenman, and Michael Charney for joining with Dr. Merigan to assist with an early evaluation of the program's performance.

Most of the computer time used for the development of MYCIN was generously contributed by Stanford Research Institute. I am grateful for the advice and assistance of Nils Nilsson, Bertram Ra-

phael, Peter Hart, Richard Waldinger, Dan Lynch, and Richard Fikes, all of whom are associated with SRI. I also want to thank Dr. Nilsson for permission to borrow his thoughts regarding the AI field's organization as discussed in Section 1.2.

In addition to Dr. Cohen, I would like to thank the other Stanford faculty who have advised me closely and who suggested both academic and research directions appropriate for my interdisciplinary interest: Prof. George Forsythe (Computer Science, deceased), Prof. Byron Brown (Biostatistics), Prof. Patrick Suppes (Philosophy), Dr. Howard Sussman (Pathology), and Prof. Cordell Green (Computer Science).

My thanks also go to the following individuals who have devoted time to useful discussions concerning the research or have assisted with manuscript review: Casimir Kulikowski, Ed Sondik, Ernest Shortliffe, Jonathan King, John Hannigan, Russ Briggs, Ken Colby, James Fries, Mark Perlroth, Roger Schank. The secretarial assistance of Claire Lustig, Sylvia Wyman, and Linda Halloran has also been indispensable throughout MYCIN's development and I thank them all for the time they have given us.

I want to close with three special acknowledgements: to Bob Greenes for his friendship and for his encouragement regarding a career combining computers and medicine, to G. Octo Barnett for his similar encouragement and for suggesting Stanford as an appropriate medical school for someone desiring to combine the two fields, and to my wife Linda, who somehow managed to adjust to our strange and often incompatible schedules, even when I and my computer terminal tied up the home phone lines for hours at a time.

This work was supported in part by the Medical Scientist Training Program under NIH Grant No. GM-01922. The computer time was provided by Stanford Research Institute under Advanced Research Projects Agency (ARPA) contract DAHCO4-72-C-0008, and by the Stanford University Artificial Intelligence Project. Secretarial and copying services were provided at Stanford University School Of Medicine under BHSRE Grants HS-00739 and HS-01544.

E. H. S.