Preface

In August of 1980, Stanford University was the site of the annual workshop on artificial intelligence in medicine (AIM). This specialized area of medical computer science research had been born almost ten years earlier with the near-simultaneous development of AIM research groups at Massachusetts Institute of Technology (in collaboration with physicians from the Tufts–New England Medical Center), the University of Pittsburgh, Rutgers University, and Stanford. These small groups of computer scientists working in the field were drawn together naturally by their common interests and by the establishment of the SUMEX-AIM network (Stanford University Medical EXperimental Computer for Artificial Intelligence in Medicine). This computing resource was established by the Biotechnology Resources Program of the NIH in 1974 and consisted of a pair of computers, one at Rutgers and one at Stanford, linked by a communications network. The funding for SUMEX-AIM not only provided computing power for researchers exploring the potential of artificial intelligence techniques in medicine but also established a series of annual workshops so that the investigators could gather to share their insight, results, and ideas regarding approaches to the difficulties they encountered.

The 1980 workshop was the first at Stanford; the five earlier sessions had been held at Rutgers University in New Jersey. Because of a growing interest in computers generally, and in artificial intelligence in particular, among local physicians and medical faculty, we decided to organize a public AIM tutorial to be held immediately following the small workshop. Most of the field's leaders were going to be there, and it seemed logical to extend their stay so that a public series of lectures could be held to acquaint a medical audience with the progress, current status, and potential of the emerging discipline. We were delighted by the interest in the program, by the excellent attendance at the two-day tutorial, and by the positive attitudes of the attendees (Teach and Shortliffe, 1981).

One of the lessons of that tutorial was the need for a readily available collection of readings to describe the first decade's work in the field. Those articles that had appeared were scattered in a number of publications, some from the medical literature and others from computer science journals. There had been no effective effort either to bring together the key articles or to describe the historical progression of work in the field.

As the word spread about the 1980 tutorial, we received increasing numbers of requests for such a collection, and the idea for this volume emerged. There was a clear need for a survey of key AIM activities, par-

ticularly in light of the success and visibility of the early efforts and because of the frequent failure to appreciate the significant barriers that remain to be overcome before widespread clinical impact will be achieved. This book is, accordingly, an attempt to address those issues. The authors and publishers of the original articles have kindly permitted us to reproduce them here, generally with only minor modifications to correct any inaccuracies that were discovered in retrospect. We have, in addition, included a new introductory chapter that defines the field of artificial intelligence, outlines its relevance to medicine, and identifies the key research issues that have guided AIM work and continue to do so. Each chapter is preceded by a brief introduction that outlines our view of its contribution to the field, the reason it was selected for inclusion in this volume, an overview of its content, and a discussion of how the work evolved after the article appeared and how it relates to other chapters in the book.

It is important to note that this book is by no means an exhaustive review of all AIM work during the period 1971–1981. Several fine pieces of work could not be included because of space limitations. We have accordingly tried to provide references to additional key articles throughout the volume. Those included were selected to provide a broad coverage of issues, as well as to exemplify what we consider to be some of the best and most influential work in the field.

The papers here tend to be more technical and detailed than those that appeared in a recent shorter volume on the subject (Szolovits, 1982). We have also provided a comprehensive index, a name index, and a bibliographical listing of all references cited throughout the volume. These additions have been designed to make the issues accessible to interested readers, particularly physicians, who may not have had previous experience with artificial intelligence. The chapters are organized in a loosely chronological way, with surveys and general system descriptions near the beginning and more recent work toward the end. The title of the volume has been selected to make it clear that we see the AIM field as a young and emerging discipline. Our views of the future, with an emphasis on the challenges as well as the promise that lies ahead, are the subject of the closing chapter.

This preface would be incomplete if we did not acknowledge and express our gratitude for the remarkable assistance we have had in preparing the book. The authors of the individual chapters dug through their archival records and provided us with copies of original manuscripts (often in electronic form) and the original figures that are reproduced in some of the chapters. Chapter introductions have been in large part adapted from the original abstracts; the authors helped us greatly by reviewing this material and correcting and augmenting the historical information (and we thank Paul Feltovich especially for providing his interesting sketch of the DIAGNOSER project).

We are also grateful to Michael Morgan of Addison-Wesley for his encouragement in bringing the volume to print, to Christopher Lane for

Stanford University   W. J. C.
May 1984   E. H. S.