Report of
Panel
4

Long Range Plan

National Library of Medicine

Medical Informatics

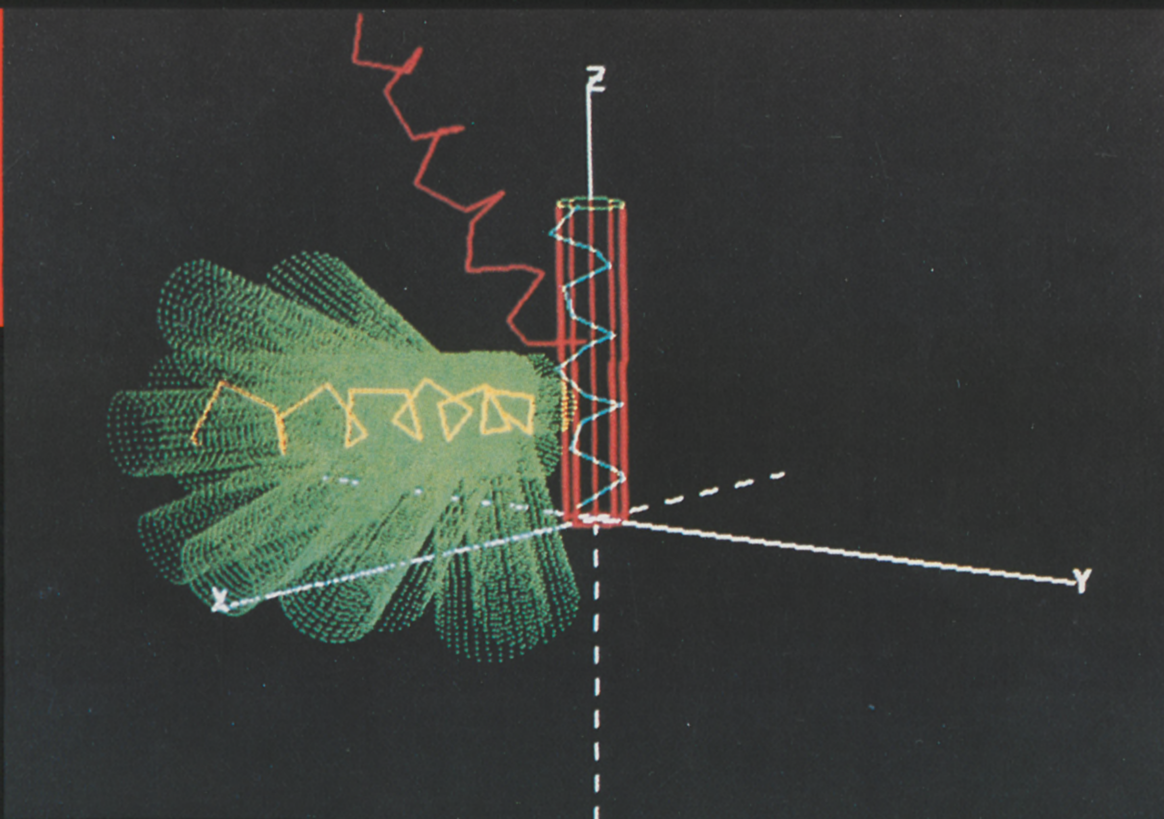

U.S. Department of Health and Human Services
Public Health Service
National Institutes of Health

Report of Panel 4

Long Range Plan

National Library of Medicine

Medical Informatics

U.S. Department of Health and Human Services
Public Health Service
National Institutes of Health

December 1986

Members and Staff of Panel 4 Medical Informatics

Chairperson

Edward H. Shortliffe, M.D., Ph.D.
Associate Professor of Medicine
and Computer Science
Department of Medicine
Stanford University Medical Center
Stanford, California

Members

J. Robert Beck, M.D.
Assistant Professor of Pathology
and Community and Family Medicine,
Director
Program in Medical
Information Science
Dartmouth-Hitchcock Medical Center
Hanover, New Hampshire

Marsden S. Blois, M.D., Ph.D.
Professor and Chairman
Section on Medical Information Science
University of California-San Francisco
San Francisco, California

Robert Braude, M.L.S., Ph.D.
Assistant Dean for Information Resources
Cornell University Medical College Library
New York, New York

Milton Corn, M.D.
Dean
School of Medicine
Georgetown University
Washington, D.C.

Arthur Elstein, Ph.D.
Professor of Health Professions Education
University of Illinois at Chicago
Chicago, Illinois

Dennis Fryback, Ph.D.
Professor of Industrial Engineering
and Preventive Medicine
University of Wisconsin
Madison, Wisconsin

Nina W. Matheson, M.L.
Director
William H. Welch Medical Library
Johns Hopkins University
School of Medicine
Baltimore, Maryland

Clement J. McDonald, M.D.
Professor of Medicine
Indiana University
School of Medicine
Indianapolis, Indiana

Judy G. Ozbolt, Ph.D., R.N.
Associate Professor
Center for Nursing Research
University of Michigan
Ann Arbor, Michigan

Ramesh Patil, Ph.D.
Assistant Professor
Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts

Stephen G. Pauker, M.D.
Associate Professor of Medicine
School of Medicine
Tufts University
Boston, Massachusetts

Thomas Rindfleisch
Director
Knowledge Systems Laboratory
Stanford University Medical Center
Stanford, California

Consultants to Panel 4

Donald A. Senhauser, M.D.
Chairman
Department of Pathology
Ohio State University
Columbus, Ohio

Homer Warner, M.D., Ph.D.
Professor and Chairman
Department of Medical Informatics
School of Medicine
University of Utah
Salt Lake City, Utah

Bonnie Webber, Ph.D.
Associate Professor
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, Pennsylvania

NLM Staff
Peter Clepper
Executive Secretary

Harold M. Schoolman, M.D.
Research Person

Earl Henderson, M.S.E.E.
Resource Person

Bruce G. Buchanan, Ph.D.
Professor of Computer Science
Stanford University
Stanford, California

Don E. Detmer, M.D.
Vice President for Health Sciences
University of Utah
Salt Lake City, Utah

Robert A. Greenes, M.D., Ph.D.
Radiologist and Director
Computer Science Division
Brigham and Women's Hospital
Boston, Massachusetts

Gwilym S. Lodwick, M.D.
Associate Radiologist
Harvard Medical School
Massachusetts General Hospital
Boston, Massachusetts

Richard Lyders
Executive Director
Houston Academy of Medicine
Texas Medical Center Library
Houston, Texas

Randolph A. Miller, M.D.
Associate Professor of Medicine
School of Medicine
University of Pittsburgh
Pittsburgh, Pennsylvania

Joyce A. Mitchell, Ph.D.
Director
Information Science Group
University of Missouri-Columbia
Columbia, Missouri

Allen Newell, Ph.D.
University Professor
Computer Science Department
Carnegie Mellon University
Pittsburgh, Pennsylvania

William S. Yamamoto, M.D.
Professor and Chairman
Department of Computer Medicine
The George Washington University Medical Center
Washington, D.C.

Members of the Panel on Medical Informatics

Chairperson
Edward H. Shortliffe, M.D., [illegible]
Associate Professor of Medicine
and Computer Science
Department of Medicine
Stanford University Medical [illegible]
Stanford, California

Members
J. Robert Beck, M.D.
Assistant Professor [illegible]
and Community and Family [illegible]
Director
Program in Medical
Information Science
Dartmouth-Hitchcock Medical Center
Hanover, New Hampshire

Marsden S. Blois, M.D., [illegible]
Professor and Chairman
Section on Medical Information [illegible]
University of California-San Francisco
San Francisco, California

Robert Braude, M.L.S., [illegible]
Assistant Dean for Information [illegible]
Cornell University Medical College
New York, New York

William Lowe, M.[illegible]
Dean
School of Medicine
Georgetown University
Washington, D.C.

Arthur Elstein, Ph.D.
Professor of Health Professions [illegible]
University of Illinois [illegible]
Chicago, Illinois

Dennis Fryback, [illegible]
Professor of Industrial Engineering
and Preventive Medicine
University of Wisconsin
Madison, Wisconsin

Nina W. Matheson, M.L.[illegible]
Director
William H. Welch Medical Library
Johns Hopkins University
School of Medicine
Baltimore, Maryland

Clement J. McDonald, M.D.
Professor of Medicine
School of Medicine
Indianapolis, Indiana

Judy G. Ozbolt, Ph.D., R.N.
Associate Professor
Center for Nursing Research
University of Michigan
Ann Arbor, Michigan

Ramesh Patil, Ph.D.
Assistant Professor
Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts

Stephen G. Pauker, M.D.
Associate Professor of Medicine
School of Medicine
Tufts University
Boston, Massachusetts

Thomas Rindfleisch
Director
Knowledge Systems Laboratory
Stanford University Medical Center
Stanford, California

Contents

Report of Panel 4

Long Range Plan

Preface

This report addresses the need for research in medical informatics and makes specific recommendations for a research program, with suggested priorities, to NLM (National Library of Medicine). There are four companion reports from other consensus panels that have made recommendations in additional areas pertinent to the activities of NLM. Although we present detailed recommendations in section five, the report first provides information to motivate the specific recommendations and explain their rationale. Section one introduces the topic, defining medical informatics and providing a general perspective on the field and its potential for benefitting both the Nation's health care and biomedical research. In section two, current and past NLM activities in medical informatics are briefly reviewed. Then, in section three, a futuristic scenario is presented, one that presents a view of how medical informatics might affect and benefit the health-care setting in two decades or so. The emphasis there is on emerging technologies and how they are likely to be integrated into familiar medical environments. The scenario is then discussed in light of the existing research challenges and barriers to progress.

Section four, the longest in this report, summarizes several key areas of medical informatics research, outlining the state of the art, existing challenges, and proposed strategies for progress towards the distant goal outlined in the scenario of section three. Then, in section five, specific recommendations are presented. These generally cut across several of the research areas from section four. They are motivated in terms of a 10-year horizon that is judged to be on a proper trajectory if the 20-year goals are to be achieved. Finally, in section six, the recommendations are summarized and prioritized. Readers wishing a brief overview of the report will find all key recommendations summarized in that chapter.

This report is the result of a coordinated effort by several consultants. Starting with organized presentations as well as open discussions, Panel members soon began to write materials that have come together in the form of this final document. Thus, the report is an effort by the full Panel, collated and edited but consisting largely of materials written by individual members. Outside consultants also contributed substantially with comments in response to earlier drafts.

It should be emphasized that the joint writing effort would have been impossible without electronic communications facilities. All Panelists had (or were given) accounts on various machines on a national network (the ARPANET), and draft text was freely exchanged and critiqued. Contributions were in many cases submitted electronically to Dr. Shortliffe and Mr. Rindfleisch at Stanford University, where files were merged and edited into a final document.

The Panelists and staff, all of whom contributed substantively to the report, are listed elsewhere in this report.

Acknowledgments

We are indebted to a number of individuals whose contributions to the planning process complemented those of the Panelists. Professor Allen Newell of Carnegie Mellon University kindly consented to help set the tone for what lay ahead. His experience in the computing field, and his warnings about the inherent folly in trying to anticipate the future, helped place in perspective the process we were about to undertake. His comments on a draft version of the report were also very helpful.

Special thanks should go to those individuals at NLM who assisted greatly in logistical support for the preparation of this document. These include Jules Aronson and Debbie Bennett, who facilitated establishment of the electronic communications network and provided training to those who were new to this technology. Dr. Henry Riecken, while trying to coordinate activities for all five planning Panels, managed to provide both moral support and useful suggestions as the final document began to take shape. Professor John Starkweather, on sabbatical at NLM, Dr. Harold Schoolman, and Mr. Earl Henderson also made several useful suggestions. Dr. Donald Lindberg's encouragement and advice often helped Panelists keep on track and focused on the specific tasks with which we had been charged. But special acknowledgment for yeoman's duty goes to Peter Clepper, whose organizational and writing skills kept the process running smoothly.

1 Background and Context

Health professionals need only pause briefly to recognize that a large percentage of their activities relate to information management—for example, obtaining and recording information about patients, consulting with colleagues, reading the scientific literature, planning diagnostic procedures, devising strategies for patient care, interpreting results of laboratory and radiologic studies, or conducting case-based and population-based research. It is society's overriding concern for patient well-being, and the resulting need for optimal decision making, that tends to set medicine apart from many other information-intensive fields. That concern gives a special significance to the effective organization and management of the huge bodies of data with which health professionals must deal. It also suggests the need for specialized approaches and for skilled scientists who are knowledgeable about both medicine and information technologies.

Medical informatics is an interdisciplinary field that combines medical science with several technologies and disciplines in the information and computer sciences. It provides methodologies by which these fields can contribute to better use of the medical knowledge base and ultimately to better medical care.[1] The emergence of medical informatics as a new discipline is due in large part to advances in computing and communications technology, an increasing awareness that the knowledge base of medicine is essentially unmanageable by traditional paper-based methods, and a growing conviction that the process of expert decision making is as important to modern biomedicine as is the fact base on which clinical decisions or research plans are made.

Each generation of scientists finds its information systems inadequate,[2] but today, not only medicine, but all sciences face a staggering information-management problem. Eighty to ninety percent of all scientists in history are alive today.[3] Their work has uncovered new knowledge in unprecedented quantities. Over the past decade, the number of publicly available data bases has grown from 200 to 3,010, 36 percent of which contain bibliographic information, and the number of records contained in data bases has grown exponentially from 52 million to 1.7 billion.[4] A decade ago, one author estimated that a core knowledge of internal medicine alone involved about a million 'facts', two million if subspecialties were included.[5]

Yet over the past 500 years, the mechanisms for storage and dissemination of this knowledge have not changed fundamentally. Consequently, the interval between discovering or uncovering knowledge and applying that knowledge to science or medicine continues to be wide because our capacity to organize and disseminate information is inadequate. In medicine there are countless examples of inadvertent failures to use life-saving knowledge because of the lack of an effective way to make pertinent information available at critical decision points.[6] For example, as late as 1972, 5 percent of physicians in one state still used a particular antibiotic to treat outpatient upper respiratory infections, even though it was known occasionally to cause fatal aplastic anemia.[7]

While most of biomedical science is directed at the study of life processes, medical informatics is concerned with the invention and dissemination of powerful information-management tools. These include, but are not limited to, (a) frameworks for organizing and encoding medical data and knowledge, (b) methods for acquiring and representing the judgmental knowledge that is acquired through medical experience rather than formal studies, (c) computer programs to permit efficient communication among health personnel, and (d) systems to provide customized advice so that practitioners may have access to expertise that otherwise might not be efficiently available when and where it is needed.

How can we make knowledge more accessible and whose responsibility is it to do so? It seems logical to turn to libraries for assistance since they have traditionally served as sources of recorded knowledge. However, until relatively recently, libraries have served as passive repositories for information, limited by the books-and-shelves approach to information storage. While researchers increasingly store experimental data in flexible electronic forms that can be analyzed, displayed, or combined with other data in many ways, libraries hold summary data frozen in traditional paper media.

The need to interrelate and combine information usefully across disciplines is increasingly recognized. A National Academy of Sciences panel recommends that "scientists start thinking about how to develop a biology-wide information system—a computerized matrix data base—structured so that it can be accessed from a multitude of dimensions."[8 9] Processing information faster or more efficiently—which today's technology can easily accomplish—is not sufficient. What is needed is more intelligent processing, logical aggregation of information, synthesis and analysis, and the development of knowledge systems that serve purposeful ends.

Basic changes in how medical knowledge and information are stored and retrieved have been called for repeatedly.[10 11 12 13] Due to work in the area of medical bibliography and medical informatics over the past two decades, it is now possible to look ahead to a true transformation of libraries and information management.

The remarkable decrease in cost and increase in power of computer hardware, plus developments in the information sciences, make a radically different kind of library possible. The library of the future will store all of its contents—images, text, data—in electronic forms. Data will include information not previously included in formal paper publications. Aggregate data from many studies, when collected within guidelines, will be used for re-analysis or to develop new hypotheses. Libraries will become gateways to disparate information sources. Librarians will provide access expertise as well as participate in research and development of new tools to issue these information sources. In time, users of these electronic libraries can expect to obtain "expert-level" advice derived through advanced computational techniques.

A health-care scenario for the year 2006, given in section three, provides a glimpse into the future by depicting how such systems might work in daily life. Parts of this scenario have taken rudimentary form even today, and the clinical and research impact of medical informatics is already being felt:

- Hospital information systems, which provide communication and information-management functions in medical institutions, are increasingly being installed.
- Physicians can search the entire pharmacopoeia in a few seconds, using the information provided by a computer program to anticipate harmful side effects or drug interactions.
- Electrocardiograms generally receive their initial analysis by computer programs, and similar techniques are being introduced for interpretation of pulmonary function tests and a variety of laboratory and radiologic abnormalities.
- Microprocessor systems routinely monitor patients and provide warnings in critical care settings such as the intensive care unit or the operating room.
- Both medical researchers and clinicians regularly use computer programs to search the medical literature, and clinical research would be severely hampered without computer-based data storage techniques and statistical analysis systems.
- Advanced decision-support tools are also beginning to emerge from research laboratories and are likely to have a profound impact on the way medicine is practiced in the future.

Other parts of the scenario are clearly possible, but substantial fundamental obstacles must be overcome to make them reality. Bringing the systems underlying the vision to a fully developed and clinically validated state will require a major commitment of resources, strong leadership, and interdisciplinary research. Progress will inevitably be dependent on high-quality, fundamental research in areas such as acquiring, representing, and flexibly reasoning about the diverse information relevant to biomedicine (e.g., instrument data, images, and direct human observations of patients). Such progress is further dependent both on the orientation and training of scientists who choose to pursue careers in medical informatics and on their assured access to advanced computing and communications resources. Detailed discussions of these research questions and opportunities follow in the section on major research issues, opportunities, and impediments.

NLM has played an important role in the early development of medical informatics. For example, its MEDLARS (Medical Literature Analysis and Retrieval System) and MEDLINE (MEDLARS Online) services broke new ground in bringing into being a new online data base industry and a national document delivery network system that affect the life of almost every biomedical scientist. Because it has been innovative in carrying out its responsibilities for the collection and dissemination of biomedical information for the entire scientific, education, and health-care community, NLM is more than a traditional library. And NLM has a special role to play over the next two decades and beyond to help insure that informatics principles are effectively developed, disseminated, and applied—both for the well-being of patients and to help biomedical professionals in their clinical and research activities.

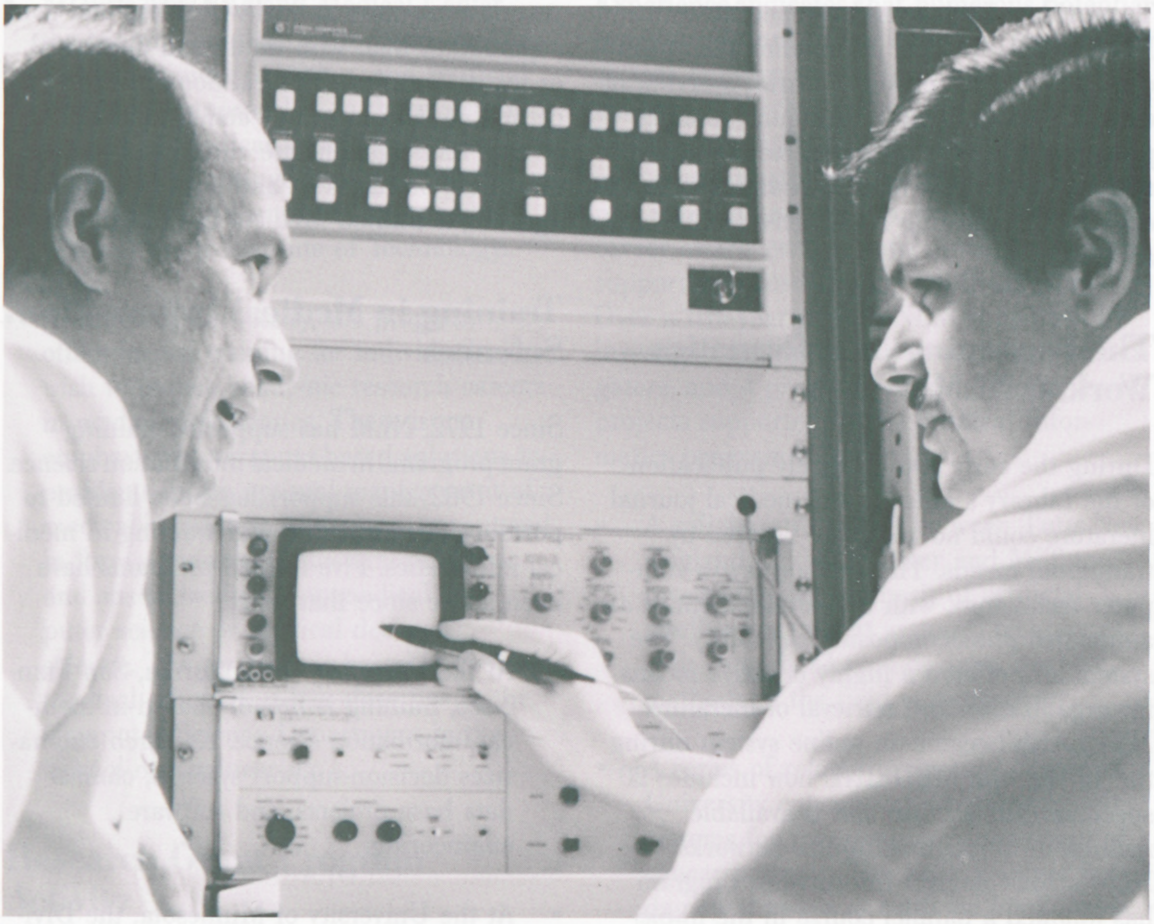

NLM's mandate is to exercise leadership in order to assure that medical information systems evolve in socially useful directions. This role extends beyond providing support for training and research programs in medical information science, and extends beyond serving as the exemplar for other libraries through its service, intramural research, training, and development programs. These roles are critically important, but the role unique to NLM is that of bringing together and mobilizing the spectrum of scientific associations, organizations, institutes, and publishing interests to undertake the task of designing and developing the biomedical knowledge base for the future. The opportunities are enormous. The barriers are not trivial.

2

NLM Programs and Recent Accomplishments

To provide a background perspective on NLM and its medical informatics activities to date, this chapter provides a brief overview of current intramural and extramural research programs. Many of the Panel recommendations in section five refer to changes in or expansion of existing NLM activities in this area.

The Beginning of Informatics Work at NLM

During the early 1960's, timely publication of the Library's index to the medical journal literature could no longer be assured by traditional means. The solution lay in computer technology, with the development of an automated photocomposing capability. The system also offered a highly effective means for organization and retrieval of literature. MEDLARS became an online system during the 1970's, and MEDLINE now includes 20 specialized data bases and is available nationwide through computer networks. NLM's Toxicological Information Program has followed a parallel course in the implementation of TOXNET, a powerful interactive online system for data on hazardous and carcinogenic substances.

Lister Hill National Center for Biomedical Communications

The Lister Hill Center, authorized in 1968, is NLM's intramural laboratory for informatics research and development. Work in this area has included the development of a computer-based compendium of information about hepatitis known as the "hepatitis knowledge base." At present, active research includes expert systems based on artificial intelligence methodologies. (Artificial intelligence is the study of symbolic reasoning techniques that permit computers to deal with ideas and concepts rather than with traditional numerical or textual entities.) Medical domains concerned are rheumatology and hematology. Several investigations are currently underway to explore obstacles to, and terminology linkages for, a Unified Medical Language System.

Training in Medical Information Sciences

Since 1972, NLM has supported training grant programs in medical information science. Since 1982, this support has been limited to postdoctoral research career training in medical informatics. Five funded programs have been active since that time:

- At the University of California, San Francisco, training is broadly based in medical information science. Research emphasizes decision-support systems, clinical data base projects, and software engineering.
- At the University of Minnesota, the Division of Health Computer Sciences gives an interdisciplinary focus to the cognitive, computer, and information sciences related to clinical management and applications for health knowledge.
- At Harvard Medical School, training is offered in major research areas of computer-based decision-support systems, representation of medical knowledge, data base and data analysis systems, and computer graphics, among others. There is a collaborative relationship with the Laboratory of Computer Science at Massachusetts General Hospital and the Harvard School of Public Health.

- At Tufts-New England Medical Center, the training emphasizes medical decision making, clinical cognition, and artificial intelligence approaches to the structure and use of medical knowledge. There is a collaborative arrangement with the Laboratory for Computer Science at the Massachusetts Institute of Technology.
- At Stanford University, the program offers formal training for individuals who wish to pursue academic research careers in medical informatics. There is an emphasis on computer-based solutions to problems in the optimal management of medical knowledge. All trainees are entered in formal masters or doctoral programs and specialize in either medical computer science or medical decision science. The program is based in the medical school, but has close ties to the Computer Science Department in the School of Engineering.

Research Grant Program Activity

NLM's extramural grant program in medical informatics is small but significant since it accounts for much of the Federal funding for research in the field. Its level of activity in recent years is summarized in the following table:

	Applications			
Year	Received	Approved	Funded	Awarded
1979	63	26	19	$1,489,453
1980	56	21	13	902,741
1981	32	16	10	733,053
1982	37	21	9	452,639
1983	26	15	10	588,456
1984	24	8	8	441,005
1985	66	24	16	2,001,588

Direct costs are shown to indicate actual resources available to investigators.

Experience indicates that numerous potential grant applicants do not apply for research awards because they believe that their chances of success are too low. By the end of 1985, however, NLM was supporting 23 active investigator-initiated project grants, 6 new investigator projects, and 6 research career awardees. The total amount of grant support for these awards was $4,297,000. Over a five-year period, NLM supported 51 investigator-initiated projects, 20 new investigators, and 4 program projects. Examples of projects supported by the Medical Informatics program include:

- A collaborative effort between Tufts-New England Medical Center and MIT to investigate the synthesis of artificial intelligence (expert systems) techniques with computer-based decision analytic methods in the field of nephrology.
- A series of related projects at Stanford University concerning advanced expert systems for therapeutic management of cancer patients. The work involves fundamental issues of artificial intelligence methodology and computer-based representation of knowledge.

- Investigation of a prototypical consultant system at the University of Missouri, Columbia for genetic maladies such as deaf-blind syndromes. This is an application of artificial intelligence to genetic diagnosis.
- Studies of computer-based clinical decision making at the Deseret Foundation (University of Utah), including radiologic diagnosis of pulmonary disease.
- The development of a major knowledge base covering all of internal medicine. This work at the University of Pittsburgh is complemented by fundamental studies of knowledge retrieval and computer-based inferencing.
- An analysis of the scientific validity of the literature of controlled clinical trials. The long-term goal of this project at Mount Sinai School of Medicine is better criteria for designing trials and editing their reports.
- An investigation of more comprehensive retrieval from bibliographic data bases. Progress at Case-Western Reserve University indicates that conventional search strategies are improved with the addition of certain attributes of semantic relevance.
- A study of image processing techniques to organize and abstract information from photomicrographs. The project at Rush Presbyterian-St. Luke's Medical Center uses images from Papanicolaou smears.

IAIMS (Integrated Academic Information Management Systems)

The concept of this program initiative is to assist academic centers with the planning and use of computer and communications technologies to create systems that integrate operational and academic information in academic health settings. NLM awards grants to institutions for IAIMS strategic planning, testing, and model development. Grants are also awarded to support research specifically related to IAIMS activities. Baylor College of Medicine, University of Cincinnati Medical Center, Harvard Medical School, the Johns Hopkins University School of Medicine, and the Oregon Health Sciences University are currently supported with planning grants.

Following successful completion of institution-wide IAIMS plans, Columbia University, Georgetown University, the University of Maryland at Baltimore, and the University of Utah are testing IAIMS concepts on a small scale. In 1985, just over $2 million was devoted to IAIMS grants of various kinds. At the grantee institutions, these grants provide a locus for numerous studies, developmental projects, and similar activities to improve information access and utilization. They contribute to organizational development appropriate for adoption of the new capabilities deriving from medical informatics research.

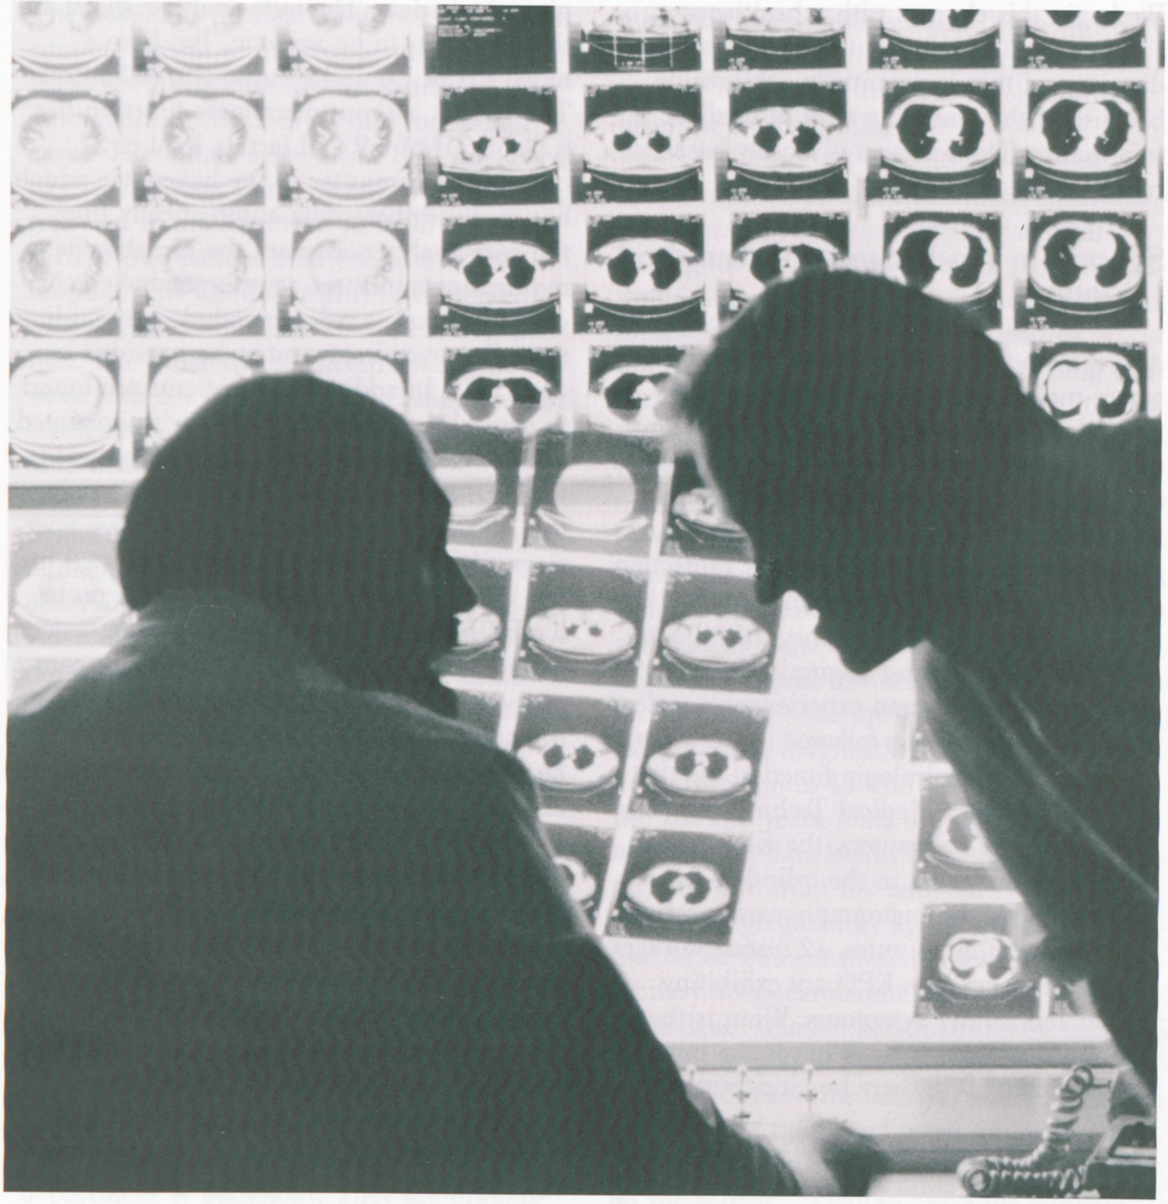

3 A Vision of the Future

We begin this chapter with a health-care scenario in 2006, a speculative glimpse of a distant goal that is potentially achieveable because of the emerging field of medical informatics. The scenario is loosely based on recent medical literature.

Scenario: An Industrial Accident in 2006

At a remote industrial plant in rural Virginia, where rocket fuel research had been performed in the 1950's and 60's, workers are detoxifying old cylinders containing unknown gases. Some gas is accidentally released, engulfing three men. The rescue squad and the company EPO (Environmental Protection Officer) are immediately summoned. By the time the air ambulances arrive, the men are gasping for breath and beginning to experience seizures. One man experiences a violent opisthotonic convulsion followed by the absence of spontaneous neurologic function. As the EMT (Emergency Medical Technicians) rush the men to the helicopters, the EPO takes samples of the gases in the cylinders for assay in a gas chromatograph/mass-spectrograph. Within 20 minutes, 12 rescue workers, 2 bystanders and the EPO are exhibiting similar but milder symptoms. What is the gas and how toxic is it? What is the immediate treatment? Will there be long-term effects?

The air ambulance medical station-data analysis unit is fully equipped for video/voice/digital data communications and analysis. While the EMTs connect the men to monitoring systems and take blood samples, the data analysis technician establishes communication links to the EPO performing the gas assay, the Toxicology Information Bank, and the receiving hospital. She verbally reports signs and symptoms and location of the accident. As she speaks, the computer simultaneously processes her words and the patients' physiologic data. The auto-analyzer shows all three patients to be acidotic. Two have a blood pH of 7.0 and bicarbonate of 5mEq/l. The patient without neurologic function has a blood pH of 6.9 and lactate level of 22 mEq/l. The computer in the helicopter, which has received these data automatically from the auto-analyzer, assesses the bicarbonate requirements in these severe metabolic acidoses and makes recommendations regarding the emergency treatment to be provided by the EMTs. In addition, all patients are found to have elevations of SGOT, brown pigmented urinary casts, and all have rhabdomyolysis, with the creatine kinase (CPK) in one rising to 47,000. The gas chromatograph-mass spectrograph assay is completed and reported to the EMTs and the receiving hospital: probable B5H9 (pentaborane), although the system recommends confirmatory studies with complementary spectral analysis techniques when the patients reach the hospital. All these data become available while the air ambulance is en route.

The EMTs find the men's personal ID wallet cards, magnetically coded like bank cards, that carry critical personal data including health information such as their medical history and baseline laboratory data. They insert the cards into a special emergency reader, which unscrambles the privacy-protection code and displays the information, including a photograph and dental x-rays for positive identification. At the same time, an admission record is automatically created at the hospital. One patient has been treated for Addison's disease; the other two haven't been to a hospital before. Immediate relatives are contacted automatically, and the EPO explains what is happening. The families arrive at the hospital shortly after the helicopters.

The hospital's medical information-decision support system recognizes pentaborane toxicity as the likely cause of the syndrome and automatically searches its files for similar cases. There have been none, but the Toxicology Information Bank identifies three cases, reported in the literature 10 to 15 years earlier. One patient died en route to the hospital. The autopsy report documented widespread central nervous system degeneration. The other two recovered within the first week with few residual effects. Several animal studies report selective reaction of pentaborane with nervous tissue.

In the ER (Emergency Room), the physician in charge and two residents have been observing the EMT crew at work via the ER video monitor and their personal workstations, which are the size and shape of the clipboards that they once would have used for note taking and analysis. Two of the patients have required blood pressure and ventilatory support. One patient has just gone into cardiac arrest. Eighteen people have been exposed; it appears that at least two will die. Because no information is available on long-term effects, the hospital medical information system establishes an individualized follow-up protocol on everyone involved.

Earlier in the day, the personal physician of one of the victims is in his office reviewing tissue specimens with a colleague from the department of pathology. They are organizing material for an article that they will coauthor. On his personal workstation, the internist examines three-dimensional electron micrographs from the digitized versions stored in the hospital's medical information system. The computer circles an area it determines is worthy of comment and awaits a dictated analysis from one of the two physicians, recognizing which one is speaking and subsequently automatically recording and indexing the dictation for the files. The other members of the research team with which they work will call up the images and listen to their colleagues' remarks before they all meet the next morning.

The screen beeps and delivers an electronic voice message from one of their clinical colleagues—a surgeon. "I'm in the OR and would like an opinion on a suspected metastatic lesion. We think the patient is unusually young to have a primary lung cancer. I'm putting the pictures through to your screen." The pathologist, who had notified the hospital system that she would be working in another office, had been easily located by the surgeon. She is an expert on lung tumors and scrolls through the frozen section slide images and chest radiographs. Before committing herself, she touches a key at the workstation; both her screen and the one in the operating suite are filled with a window summarizing diagnostic data on all women under the age of 45 who had lung lesions in the last two years. Data are displayed from her hospital data bank, from the state tumor registry, and from the national end-results registry, segregated by age groups. "It looks like an 85 percent probability of primary pulmonary adenocarcinoma. Let's look at outcome data." Another keystroke and the under-45 group is displayed by treatment and two-year survival, disease-free survival and a stochastic projection of quality-adjusted life expectancy. The most successful protocol is highlighted on the screen. "Thanks," says the surgeon. "I think we'll go ahead."

The pathologist leaves for an appointment elsewhere, and the internist turns to a book chapter on which he is working. He is behind, but so are a couple of the authors of the other 20 chapters who collaboratively maintain this online textbook. It doesn't matter as much now as in 1986 when all 20 authors had to get their manuscripts together for several editing iterations before going to press. Now, as he revises his chapter, it automatically goes through a peer review, an editorial process, back to him, and then to online publication. Although physicians can order paper copies, computer tapes, or compact disk versions (depending on their preferences and resources), by far the easiest way is to scan the latest online versions for new information. Some people still like and need the portability of paper. Of course, the new data are also included automatically in specialized data banks and knowledge bases for decision support. His colleague from pathology couldn't have been so sure of the lung lesion otherwise. Sections like his are useful for continuing education and refresher reading.

The internist—an academician—spends a day each week on his scholarly work. He writes a book chapter, contributes to one knowledge base, and reviews an expert/decision support system. He likes to revise his chapter every three months, at least. He also uses citation mapping techniques to assess the impact of his work. He sees that the hospital library has automatically stored in his files about 20 bibliographic citations and abstracts. A great deal of information is still disseminated in the form of articles, but they are shorter, more analytic and more interesting to read, now that the data themselves are managed through networks of knowledge banks. Publishers license libraries to manage the long-term storage and short-term dissemination services. NLM, now also known as the National Institute for Medical Informatics, along with the various professional associations, coordinates the networks. Back in the early 1980's, it led the way to dynamic online textbooks with the original hepatitis knowledge base.

Citation mapping, a method of statistically relating the articles cited in research papers, is also a good way to survey a whole field and spot emerging new related research areas. Many physicians check these as well as the knowledge base to which they contribute. They often like to scroll through all the articles quickly and then check through the abstracts, which can be displayed ranked by relatedness to a given topic, to segments of the knowledge base, or a given patient, through an elegant automatic process. The physician works with a split screen: his text on the left and the new material on the right. He uses a hand-held screen outliner to cut and paste copy. Rewriting will be much easier when he can just use eye-scanning to move lines around the screen, but that is probably a few more years down the technology path.

When the author finishes, he uses a simple voice command to send copies to the reviewers and editor. His publisher maintains the peer review roster, and his reviewers change every three months. Depending on the style of the reviewer, he gets voice mail comments or written comments arrayed next to the relevant text. The editor makes her changes in electronic magenta for easier identification.

Today, our young internist has just enough time to try out a new expert DSS (decision support sector). He enjoys this activity because it tests his knowledge and keeps fine-tuning his specialty DSS. He is an author of one of the knowledge sectors and a regular critic of two others. It is a lot of work, but is now recognized as having as much academic value as writing articles or book chapters. He expects to be promoted to full professor next year. After his first test case gives surprising results, he decides to examine the logic and the knowledge base. He sends a voice mail message suggesting a review of one segment of the logic. Most of the time he uses electronic voice mail, but occasionally he takes part in the monthly conjoint reviews at the hospital, which are convivial and stimulating. Before he has a chance to start the second case, his monitor beeps. It is an urgent message from the hospital.

When he answers, he sees on the screen that his patient with Addison's disease has just been admitted to General Memorial in deep coma, following an accident at the plant exposing him to pentaborane. The cumulated data and the emergency room physician's report cascade across his screen. Two of the other twelve who have been exposed are also his patients. He leaves for the emergency room immediately.

Four months after the disaster, the internist assesses the damage. One patient—the man who initially lost neurologic function—never recovered consciousness and died on the eighth hospital day. His patient with Addison's disease is quadriplegic, blind, partially deaf, and no longer sentient. The third man was luckier: he had been the farthest away and had been able to cover his face and hold his breath. Twelve weeks after the accident, he and seven others who were exposed still evidence mild brain dysfunction and psychiatric symptoms. Elevated CNS neurotransmitter levels and abnormal ventricular brain ratios on CT scan indicate neurotoxic damage. None of this is consistent with earlier data in the Toxicology Data Bank. This isn't surprising since post trauma monitoring was not as easy and sophisticated 10 years ago when the previously reported accident had occurred.

The internist, the emergency room physician, and a nurse epidemiologist have persuaded most of those involved in the accident to participate in a long-term follow-up study. Except for coming to the hospital one day a year for a physical examination and biochemical assays, other data (such as psychological performance) are collected over the phone after voice prints have been made. The patients will be interviewed by a computer program, which will be polite and able to answer their questions satisfactorily.

The nurse-epidemiologist makes two home visits as a part of his routine follow-up protocol. His clipboard workstation is equipped with a microrecorder that tapes all interviews. A pressure-sensitive screen allows him to follow the protocol and code the responses easily. Later, the stored responses will be "uploaded" to a machine "trained" to accept both voice and digital input. A series of programs presort the information for later review by the research team. Preparing the information for publication and eventual inclusion in the Toxicology Information Bank is easy. Making the home visits complicates his normally full schedule. He usually makes between five and eight such visits in a morning, thanks to the hospital scheduling software that not only organizes the visits for optimum use of travel time, but times them according to the nature of the problem, the previous visits, if any, and the personal schedules of each patient. Driving between visits, he listens to a review of the cases and prepares to make the most of the visit. Although the next patient he will visit had been lucky with the pentaborane (sustaining only minimal brain dysfunction), within a month after the accident he has had to undergo a bowel resection for cancer of the colon. His wife has not slept well since the accident and is depressed and anxious.

The couple are relieved to see the nurse. Neither was confident about the details of colostomy care. He shows them how to use a small computer, about the size of a book, that he plans to leave with them. It "reads" its program from a compact digital video disk that includes moving pictures of proper techniques. It is programmed so that every day it lists what needs to be done. When they finish a procedure, all they need to do is run a finger across the instruction. If they forget, the next day the instruction will flash. And if they skip a procedure more than once, the computer will alert a visiting nurse to give them a call. If nothing is done for a day, an emergency alert will contact, by an integral modem, the hospital medical information system. If the couple needs something explained, all they have to do is activate the "Help Panel." The computer can distinguish between an urgent need for help and a reminder or information need. If the need is urgent, the call is immediately referred to a nurse or a physician. Otherwise, the appropriate instruction is displayed on the screen. The instructions are resident in the computer disk memory. The computer can also be used to contact their physician and to renew prescriptions. There is always a graded set of help levels easily within reach.

As the nurse-epidemiologist drives back to the hospital and the research team, he reflects on his morning. He wasn't just caring for people; he was also collecting data and using technology to enhance the selection of effective interventions and outcomes. The data from his community are automatically merged with those collected in other states, producing the best information available for patient care. He remembers the regret and guilt that he used to feel as a young professional 20 years ago when he knew he didn't have time to read and assimilate even a fraction of what he needed, much less contribute to the literature. The knowledge base in nursing and medicine is in constant flux, but now there is a system that helps keep information flowing from the patient through the health-care management team to the researchers and policy analysts and back through the loop again in an orderly, useful fashion.

There will be long-term follow-up on the pentaborane incident and the nurse-epidemiologist will see more of the couple for some time to come. But the research team is nearly finished with the first draft of their report. The data are compiled and available in the Toxicology Information Bank, labelled as preliminary because additional data are still coming in. Toxic spills are still not commonplace, but far too many toxic chemicals were buried years ago to be confident that the data won't be needed again.

Assessing the Future: How Do We Get There?

It is natural for the reader of the preceding, perhaps optimistic, view of the future to ask whether the scenario outlined there is feasible, how much of it is simply wishful thinking, and in what time frame we can expect the innovations to occur. Although the framework of this report is an attempt to look 20 years into the future, some aspects of the scenario described are many years farther away, whereas some can probably be achieved sooner.

Part of the challenge for the medical informatics field is to identify realistic matches between what ought to be undertaken (because of obvious societal good or response to a clear need) and what is appropriate given current scientific knowledge, opportunities for progress, impediments to achieving the goals, and the value of the idea relative to the costs incurred for its development or maintenance. In the sections that follow, we begin by describing the opportunities that exist for the field today—opportunities based on state-of-the-art hardware and software accomplishments, as well as logistical issues and the perceived needs of the intended users of medical information systems.

We conclude by summarizing barriers to progress. As we shall describe, these tend not to be in the area of hardware development but rather, to be largely logistical (e.g., personnel development, integration of systems into research and health-care settings, and technology transfer) and in the area of system design and basic computer science research. The future of medical informatics is not limited currently by the power, size, or cost of computers.

After this introduction to the issues, section four provides detailed discussions of several major research areas or enabling activities that need to be undertaken if even part of the 20-year scenario is to be achieved. A summary of key recommendations and some suggestions for resource needs and relative priorities are proposed in sections five and six.

Opportunities That Can Guide the Field

The opportunities for medical informatics derive from several sources, including the rapid growth of awareness and interest in computers and information-management systems by health professionals of all types, the growing usefulness of medical information systems for helping with manifest biomedical research and health-care problems, and the continuing rapid development of the technological base of these systems so that they can be made increasingly useful. Only a decade ago, computer and software tools were largely unused by health personnel and were viewed with some skepticism. The change in attitudes has been due in large part to three influences: (1) the emergence of microcomputers and easy-to-use software, with the concomitant demystification of computers and a general sense in society that such machines are somehow manageable and useful; (2) the growing distress among health professionals regarding the amount of information they need in order to practice good medicine; and (3) increasing recent fiscal pressure that encourages the practice of cost-effective medicine, thereby forcing practitioners to consider carefully the clinical utility and reliability of tests, procedures, and therapies—especially when they are expensive or risky. As a result of these influences, professional societies in medicine, nursing, pharmacy, and other allied health fields are attempting to respond to a general demand for information and education in medical informatics.

NLM has traditionally directed its attention to health-care providers, offering bibliographic services and supporting research intended to enhance the delivery of information and knowledge to a wide range of medical professionals. However, the effective and accepted systems of today (for example, MEDLINE, or clinical information systems in hospitals or clinics) may be viewed as directly or indirectly benefiting other members of the biomedical community—patients, public health workers, clinical and basic science researchers, patients, and those involved in preventive care—in addition to health-care providers themselves. As the technology advances in the years ahead, with the data bases of today progressing to online texts and journals, comprehensive knowledge bases, and ultimately integrated advice systems that directly help with decision making, NLM will continue to serve a broad community of scientists, practitioners, patients, and other health professionals.

In response to a new awareness of the potential clinical role of computers, professional societies and educational organizations have begun to call for focused programs in the area and have identified NLM as the logical Federal entity to coordinate medical informatics activities. The recent AAMC (Association of American Medical Colleges) report on the *General Professional Education of Physicians*, for example, explicitly called for increased use of informatics techniques in the medical curriculum, suggested that health professionals should routinely receive limited training about computers and their use, and recommended the establishment of academic units of medical informatics in all medical schools. Students have begun to demand such training, and continuing education programs on the subject of computers and clinical computing have become popular with practitioners. Articles on electronic publishing have appeared both in the lay press and in medical journals; they have helped create a demand for more information about how computer technology will affect the way that clinicians and biomedical researchers will someday access the information that they need for their work. The routine use of computers in hospitals, and the promise of NLM's IAIMS program have also led to an increased awareness of the clinical role of computers and NLM's activities in the area. Furthermore, as individual academic institutions have begun to grapple with issues of instructional and research computing, administrative data processing, information dissemination, and communication within the medical center, the need for expertise in the area of medical informatics has become increasingly evident to administrators.

Health professionals have begun to identify professional problems with which computers could assist, but they have felt frustration at the lack of available advice about what is feasible, what has already been done, and how to avoid duplicating both the errors and the successful work of others. Professional societies have started to ask whether NLM could play a role in developing directories of available software and in disseminating such information.

Concerns about spiraling health-care costs have resulted in new practice models and reimbursement schemes in which there is pressure both for accurate diagnosis and reduced costs. Thus, there is a new opportunity to provide tools for health-care decision makers, ones that provide information about diagnosis and management but that are also sensitive to issues of efficient use of time and avoidance of excessive expenses. Decision-support technologies, including large data base tools and expert systems, have matured greatly in the last decade, and the medical informatics field is poised to begin to make such decision aids available if logistical issues can be overcome (see below). Sensitivity to 'human factors issues' and the use of novel graphics and interactive techniques have led to the development of prototype systems that are much more likely to be useful to health professionals who are unsophisticated in the use of computing technologies.

There are other developing technologies providing new opportunities for the medical informatics field and assisting in achieving the goals for the future outlined above. These include supercomputers, with their remarkable capacity and speed, and new microprocessor techniques that are placing the computing power of a large million dollar computer from 1970 on a thumbnail-sized chip selling today for a few hundred dollars. The emergence of small portable computers has suddenly made it practical to imagine placing advanced computing power in the hands of each health professional in a hospital, clinic, or office. Networking technologies will allow individual machines to share information and work in unison when necessary, while permitting independence when untethered processing is more appropriate. Progress in speech understanding and speech generation suggests that systems will be able to converse with users in the years ahead. At the same time, new graphics techniques, including high-density displays for even the least expensive personal machines, are beginning to change the way we think about communicating with computers. The diagrams and drawings typical of a patient-care chart or journal article may soon be equally well handled by the machine on the desktop.

Mass storage techniques, with methods for distributed data processing, are making feasible the notion of large, nationwide health data bases. With ever-improving techniques for maintaining data privacy, integrity, and confidentiality, it has become practical to mount efforts to standardize data formats so that health assessment, planning, and focused clinical trials can be greatly facilitated. There are major challenges to establishing such large data bases, not the least of which is the lack of a Unified Medical Language System that could assure some uniformity of information gathered at multiple institutions across the country.

As this last example has demonstrated, new technologies are providing opportunities that, in turn, highlight some of the major impediments to achieving the distant goal. The next section describes additional barriers to progress that must be recognized and addressed in any practical plan for the future.

Impediments to Progress in Medical Informatics

Although individual research areas have been restrained more by some issues than others, there are several recurring problems that cut across essentially all medical informatics research activities. Specific issues are addressed with the individual research topics in section four, but this introductory discussion summarizes the recurring themes that broadly affect the entire field.

Perhaps the greatest impediment to progress in medical informatics has been its inherent newness and its failure to be recognized as a powerful base on which medicine can and should draw. Many observers view the field as an engineering discipline and fail to recognize that many of the research issues to be addressed are fundamental. This misconception tends to lead to inappropriate expectations about the pace at which progress in the field can occur, and about the levels of support needed for progress. Because the inherent difficulty in the research is often misunderstood, it has been difficult for the field to compete effectively for the limited resources available to support medical investigation. Despite the small size of NLM's total budget and its commitment to support bibliographic and data base services, the Library is responsible for much of the research activity in medical informatics. Other NIH agencies have no direct programmatic interest in medical informatics, although categorical institutes do occasionally support specific projects and the NIH Division of Research Resources has provided advanced computing resources for some members of the biomedical research community.

Equally problematic is the dearth of individuals trained at the interface between informatics and biomedicine. NLM has used some of its limited resources to support training in the area, but the training time is generally long, and the number of graduates inadequate for the need that exists. As health science centers begin to respond to the need for academic units in medical informatics, they are finding that it is extremely difficult to identify candidates with the training or experience appropriate for an academic career in the field. In the past, there was uncertainty about the job opportunities that would be available for individuals with training in medical informatics; this concern is gradually resolving as academic institutions begin to seek researchers who can do investigative work while teaching medical informatics and assisting in the development of information-management systems for the health science center.

It must be emphasized, however, that academic health science centers have traditionally been slow to change, especially when they fear runaway resource requirements—a common issue where medical computing in general and academic medical informatics in particular are concerned. The cautious interest of institutions and health professionals is also tempered with skepticism, a modicum of mistrust in the technology, and traditional conservatism and rigidity. Although academic units and centers of excellence are beginning to emerge, the changes are likely to be slow without active encouragement and the development of improved external funding mechanisms.

There are also scientific issues that serve as impediments to progress in medical informatics research, including the need for a clearer definition of the evolving discipline itself. There are many fundamental research issues that remain, and these tend to require resources and recognition similar to those of basic science activities in the traditional medical sciences. Many of these problems are outlined in the research discussions of section four below. An overriding concern, for example, is the need for a Unified Medical Language System that will provide standards for the communication, indexing, retrieval, and structuring of data and information.

Finally, the proper methods for transferring medical informatics technology from research environments to clinical settings are currently poorly developed. Except for the emergence of data base purveyors, vendors of hospital information systems, and office management companies, the medical informatics industry is essentially nonexistent. The lack of well-defined technology transfer methods and the tendency for research prototypes to transfer poorly to new clinical environments have resulted in serious impediments to the dissemination of systems and techniques. Uncertainties regarding legal liability have also discouraged companies from developing or distributing decision-support tools. Further confounding the problem is complexity in defining the role of NLM in facilitating technology dissemination or providing clinical services; this is particularly the case when observers view NLM activities in the area as being potentially competitive with the private sector.

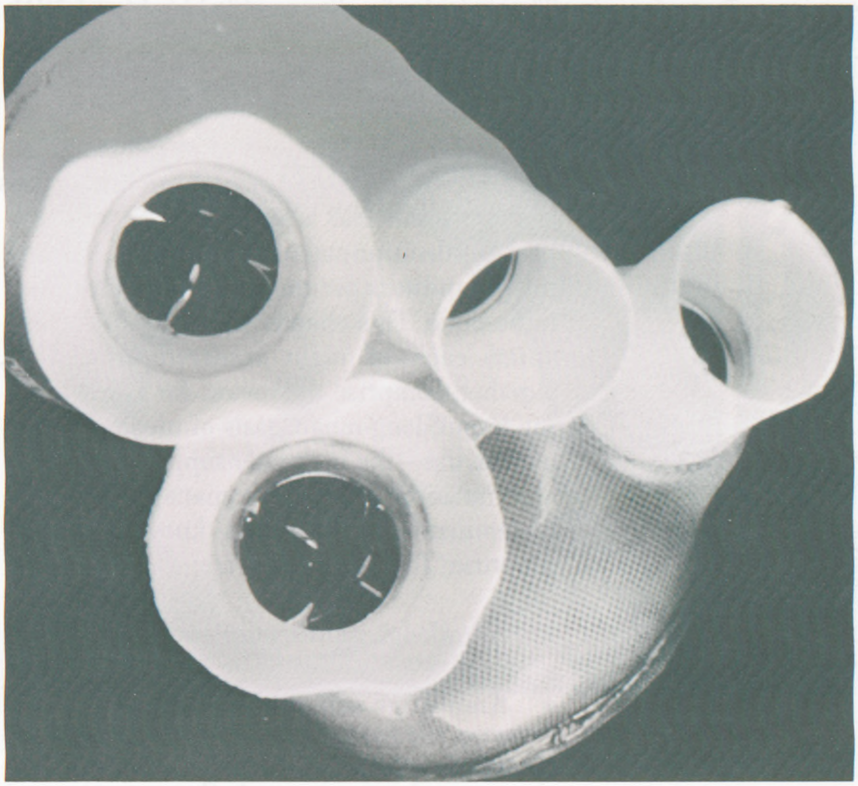

An artificial heart.

Despite the problems cited, the accomplishments and potential of medical informatics research are real and are clearly directed at the distant goal described earlier in this chapter. In the next chapter, the Panel provides a summary of key research areas in medical informatics, describing the state of the art and making recommendations for further research.

4 Major Research Issues, Opportunities, and Impediments

Through the funding and leadership of NIH (National Institutes of Health) and NLM, a community of researchers has been active for over two decades in wide-ranging applications of computers to medicine. This research has greatly advanced the capabilities of the computer as a research tool and medical decision-support system and has laid the groundwork for the substantial work yet to be done. This chapter discusses in some detail each of eight research areas that, together, comprise the means for reaching the broad, long-range goals of medical informatics—the computer-supported accumulation, structuring, management, and dissemination of biomedical knowledge and expertise. The eight areas are:

- Knowledge representation
- Knowledge and data acquisition
- Medical decision making
- Cognitive aspects of decision support
- The human-machine interface
- Information storage and retrieval
- Technology transfer and dissemination
- Supporting technologies and enabling activities

The descriptions of the respective research areas include summaries of the rationales for the research, current status, impediments, goals, and optimal strategies for achieving the identified goals.

Knowledge Representation

Rationale for research in the area

The representation of medical knowledge—its expression in mathematical and symbolic form—has always been critical to the processes of learning, knowing, practicing, and teaching in the health professions. Until recently, representation has been limited to the traditional tools offered by natural languages, two-dimensional graphics or diagrams, and photographic images. Information must be represented in permanent form in order to store it conveniently, transmit it to others, and process it automatically.

More recently, information technology has begun to bring us powerful means to assist with these tasks. The increasing availability and use of information systems have created new opportunities and raised new problems in representing medical information in forms suitable for automatic processing. Rich representation systems are central to the effective use of knowledge in medical informatics, and these must take into account the diverse forms of biomedical information and knowledge. They must capture the intricate semantic structures and interrelationships to facilitate the incorporation and synthesis of the steady flow of new information and the retrieval and use of existing information relevant to particular research and clinical support needs.

Recent accomplishments

Early in the development of computer science, means were developed for representing the numerical values of data and other non-numeric forms of information, such as text, in electronically processible forms. More recently, in the branch of computer science known as artificial intelligence, the focus has been on methodologies for representing symbolic knowledge of various kinds—factual knowledge, experiential knowledge, judgmental knowledge, or problem-solving knowledge. A number of useful techniques have been developed. Knowledge representation schemes are used to represent inferential relationships as well as properties of objects and relationships among concepts. These tools make it possible to combine multiple pieces of information, to reach decisions and give advice, or to construct processes that resemble human cognitive behavior. Examples include biomedical systems like DENDRAL for determining molecular structures,[14] MYCIN for diagnosing infectious diseases,[15] or INTERNIST-1 for diagnosing problems in general internal medicine.[16]

The study of knowledge representation has been a central problem of the field of artificial intelligence since its inception. Over the past 25 years, a number of approaches have been explored.[17] [18] These generally fall into two classes: The first makes use of logic (e.g., rule-based or production systems, or first-order predicate calculus), with well-developed semantic theory and rules for inference. The second uses semantic networks and frame-based systems with well-developed methods for structuring knowledge about the properties and relationships between objects, facts, concepts, constraints, and other entities relevant to symbolic reasoning systems. Over the last few years, systems have begun to emerge that combine the strengths of both these approaches. Significant and rapid progress in these areas is anticipated over the next several years.

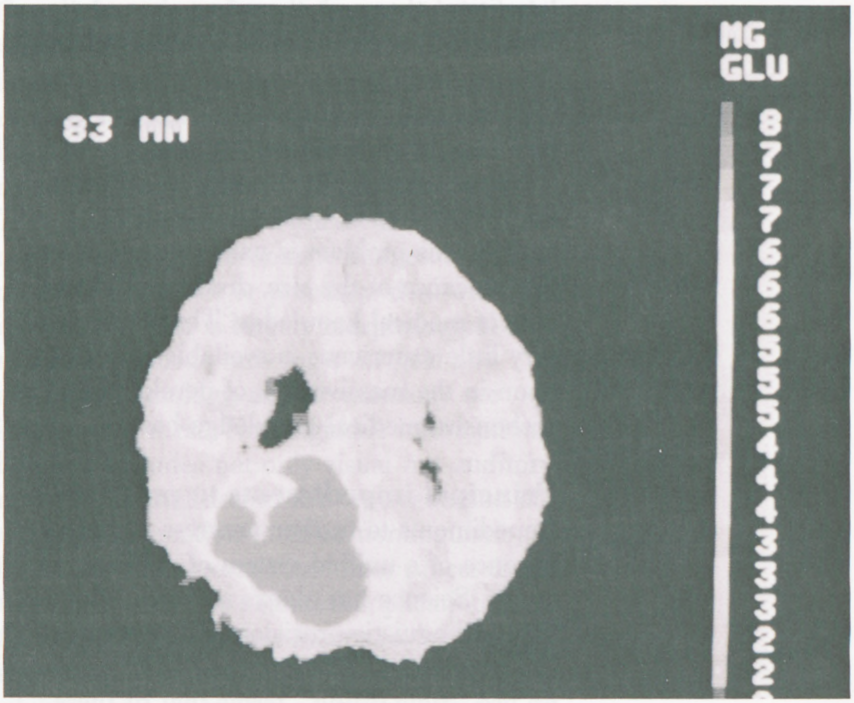

PET (Positron Emission Tomography) image.

In spite of these advances in representation formalisms, and their widespread use in "expert systems," only a few attempts have been made at building very large and complex knowledge bases such as are routinely encountered in biomedicine (see, for example, the INTERNIST-1 system[16] and the CYC system).[19] This problem is particularly relevant here because of the size, diversity, and richness of medical knowledge. Thus, at present, very little experience is available to draw upon in the massive task of developing comprehensive medical knowledge bases.

Principal impediments to progress

Impediments to further progress include the absence of a unified system of medical language. Science has always required a high level of linguistic precision. When we refer to the attributes or values of the "hard" sciences, we use "tight-fitting" terms that fit these concepts exactly. In the "softer" sciences, which include much of biology and medicine, concepts deal with more variable phenomena. Hence, we need terms with a somewhat looser fit in order to insure that we grasp the concept of interest. Clear explanations require appropriate and consistent use of language. More recently, the need to manage large volumes of medical information by machine (including patient data, bibliographic information, and medical descriptions of various kinds) requires that this be done reliably and economically. Unlike humans who can tolerate ambiguity and make powerful use of metaphors, computers require uniform language standards. NLM has already indicated that the development of a unified system of medical language will be important to its future progress. This goal merits high priority.

Knowledge appears to have many features that can be regarded as generic and thus independent of particular subjects. Were this not so, the field of cognitive psychology, for example, could have no general goals. Hence, it is common practice to speak of high-level knowledge, to invoke the notion of levels of abstraction, and to consider processes through which raw observations—data and facts—can be acquired in particular contexts and subsequently emerge as information. It also makes sense to regard information as something that, after further processing, becomes converted into knowledge. These concepts of data, information, and knowledge can be articulated without reference to particular examples of scientific disciplines. Research offers the opportunity to explore more thoroughly what these different terms and processes imply in the context of medical informatics.

There is substantial evidence that the study of knowledge is not only a branch of philosophy, but a field rich with experimental opportunities for examining the creation and recording of knowledge in intelligent systems. For example, it has been proposed that scientific knowledge may have structural properties that depend upon the hierarchical level of the descriptions involved, and that this structure accounts for the systematic differences found between the descriptions employed in, for example, physics and biology. A National Academy of Sciences committee recently proposed that biological knowledge can be represented in a matrix form by combining these (horizontal) hierarchical levels with the vertical taxonomic classes of particular biological objects.[20] NLM should give high priority to these analytical studies of the state and nature of knowledge and its interrelationships, so that automatic methods may be developed for classifying, storing, and retrieving health information. A deeper insight into this analytic process would greatly increase our understanding of the knowledge integration process and would help in the development of new tools to use in managing information.

Along with the study of medical knowledge and the development of a uniform language system for expressing this knowledge, we must also concern ourselves with representing knowledge in computer-usable forms. In other words, we must deal with not only the terms and sentences used to describe medical knowledge, but also with the underlying meaning of the concepts being used. Furthermore, these concepts and relations must be organized and indexed to provide intelligent and efficient access to knowledge and support the mechanized reasoning necessary in building knowledge-management systems that deal with various aspects of the health-care professions.

Natural language text is, of course, the most widely used means of representing medical information and communicating medical concepts and knowledge. It can be thought of as lying at one end of a continuum, where humans deal with information most naturally and where computers do so with great difficulty. At the other end, we have the opposite situation; information or knowledge is represented in a form suitable for computers but relatively inscrutable to humans. We may think of constructing information systems positioned midway along this continuum, using knowledge-representation forms that are less than optimal for both man and machine. Alternatively, we could allow both humans and computers to function in their preferred states and attempt to develop technology that will permit them to communicate with each other, translating between their respective representation systems. The impediments to this goal continue to be those problems common to computational linguistics generally: ambiguity, reference, and pragmatics (in the linguistics sense).

Certain branches of medicine are particularly suitable for representing information in nonlinguistic forms. Indeed, some of the most exciting recent developments in medicine have been means for producing images of medical objects such as CT (Computerized Tomography) scans and ultrasound and nuclear magnetic resonance images. Some medical specialties or processes deal with information that is largely morphological in nature; diagnostic radiology, surgical and anatomic pathology, and dermatology are examples of these. Nonlinguistic information, however, includes not only visual but auditory material. In addition to the representation of medical information in these nonlinguistic forms (for their storage and later processing), there is an exciting potential application in the use of nonlinguistic means for accessing stored information. This may offer novel possibilities as a means for accessing clinical data bases.

Summary of research needs

The topics discussed as impediments represent fundamental issues in medical informatics that require substantial basic and applied research. Because they are concerned with the fundamentals of medical informatics, the research issues are largely independent of the specific types of medical knowledge involved. Increasing our understanding of medical informatics at this basic level will have a great multiplying effect on the progress of the field and, accordingly, on the emergence of novel computer applications in health care.

Consequently, research goals in this area are:

- To gain an improved understanding of the structure, properties, content, and usage of medical language;
- To determine the specific requirements for and the means of developing a standardized system for encoding medical language;
- To understand better the nature and structure of data, information, and knowledge, the relationships among them, and optimal means for representing these in machine-processible form; and
- To understand better the representation of information that involves processes (e.g., information involving time or cause and effect), and the representation of nonlinguistic (nontext) information.

Strategies for future research

The difficulty and deep implications of knowledge-representation research for medical informatics make it a central focus for research support over an extended period. The preferred strategies are:

(1) Fund as many promising research efforts as possible to investigate the numerous fundamental issues surrounding the representation of biomedical knowledge.

(2) Energetically pursue the definition and adoption of a comprehensive unified system of medical terminology and language.

(3) Fund experimental efforts to encode portions of existing biomedical knowledge bases into prototypical representation systems resulting from basic research. This should include studies of how to automate this process.

Knowledge and Data Acquisition

Rationale for research

In recent years, there has been rapid growth in the generation and accumulation of information. Accordingly, much old information and knowledge quickly become obsolete with advances in the understanding of fundamental biological processes and changes in technological, environmental, and social settings. To function adequately in this changing environment, health-care professionals and researchers must spend a significant portion of their time and energy acquiring and digesting information.

Because members of the biomedical community generally lack resources to handle the large volumes of information being produced, much of the detailed data gathered during research and clinical studies is not published with the research report and is essentially lost. The validity of the reported results cannot be independently verified by other researchers because they cannot access the original data. Moreover, research hypotheses that were not anticipated during the design of the study cannot later be tested without a fresh collection of data, usually at significant cost to the researcher and funding agencies.

These problems can be significantly alleviated if the process of acquiring and assimilating information into coherent and electronically accessible sources of knowledge can be automated and incorporated into the library of tomorrow. Unlike libraries of today, which serve as passive sources of information where librarians guide users to physical documents that may contain relevant information, libraries of tomorrow, through the use of decision-support and knowledge-based technology, will be able to ascertain the specific needs of their users, extract and assimilate relevant information from their knowledge/information banks, and present it in a coherent, interactive manner. Unlike today's geographically centralized libraries, libraries of tomorrow will be accessible at the touch of a button, at home, or in clinics or research laboratories.

Such an approach will allow libraries to provide additional services that will increase their impact on the productivity of research efforts and on the delivery of quality health care. For example, the availability of large and comprehensive knowledge bases will allow libraries to provide electronically decision-support services to health-care professionals, thereby improving the quality of health care in situations where human expert opinion is not otherwise available in a timely fashion. To realize these goals, however, the processes of building large knowledge bases and providing tools to update and manage them must be key foci for medical informatics research.

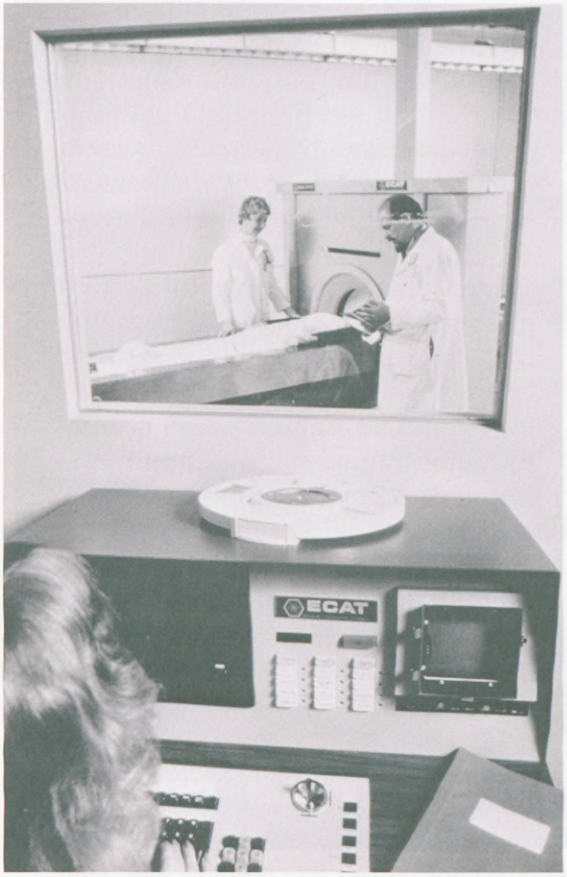

Recent accomplishments

Recent accomplishments include the application of computers to assist in clinical decision making. Over the years, many prototype systems have been developed that can assist clinicians in diagnosis,[21 22] critique approaches to patient management by prompting if the proposed action contradicts accepted medical practice,[23] and assist in management of difficult cases.[24 25 26] Large patient data bases have been developed to assist with selecting optimal therapies.[27] Methods are being investigated to extract causal relations from such data bases.[28] More recently, electronic textbooks that allow clinicians to obtain diagnostic information tailored to case-specific diagnostic situations are also being developed.[29] These experimental systems have clearly demonstrated the feasibility and usefulness of such an approach. Substantially more research and development must be done, however, before these systems can be widely used and accepted.

A key to the development of these systems has been the acquisition of pertinent information and the creation of the knowledge bases required to implement these ideas on a larger scale. In most cases, the required knowledge bases have been manually constructed, leading to what is commonly known as the "knowledge acquisition bottleneck." Recent accomplishments represent the first steps in automating this process through interactive transfer of expertise,[30 31] automatic indexing, and reorganization of knowledge bases. Text understanding promises to speed up significantly the process of knowledge acquisition and make it possible to develop large and comprehensive knowledge bases.[32]

A system for quantitative DNA analysis using image processing.

Principal impediments to progress

The cognitive process used by human experts in organizing, assimilating, and using information in diverse forms, such as raw data from laboratory and clinical studies, individual research publications, review articles, specialized monographs, and textbooks, must be better understood before significant progress can be made in building computer assistance for this process. Significant further processing is needed to go beyond simple search and retrieval of raw data and textual information. The information acquired from different sources—clinical experiments and research reports—must be integrated and generalized. Results must be reformulated and conclusions drawn that are useful to the medical community.

Although considerable progress has been made in the areas of automatic indexing, learning, and drawing knowledge from examples,[33] the technology is still embryonic. Most of the knowledge bases in use have been constructed by hand, in an extremely tedious and time-consuming process. For example, the INTERNIST-1 knowledge base has already consumed approximately 30 person-years of work from researchers with expertise in both the knowledge-base methodology and the domain of application.[34] Few such experts are available and their training is time consuming.

Much of the diverse information used by biomedical scientists is collected and prepared for publication in electronic form, but it is made available to the community at large only in print. Thus, a key impediment in the use of computers to accumulate and manage this information has been the need to convert this information back to electronic form.

Even if data and publications were directly available in machine-readable form, the diversity of software packages, operating systems, communication protocols, and formats used for storage and communication of this information would still make wide acquisition of this information a difficult problem.

Finally, managing this vast collection of information is a mammoth task. The information acquired must be indexed and stored in a manner that lends itself to efficient search and retrieval. Care must be taken to ensure the integrity and consistency of information drawn from various sources. A lack of uniform terminology and language for describing medical knowledge makes it difficult to perform these tasks effectively.

Summary of research needs

Research needs include developing new methodologies to create, accumulate, integrate, validate, and search large data and knowledge bases. This process is of central importance to the Library's ability to expand its effectiveness as the source of information and knowledge and to provide expert assistance in the retrieval and use of knowledge, simulating more and more of the human cognitive processes involved. The research needs in this area can be divided into three classes. First, community-wide standards must be developed that will allow the Library to acquire, integrate, and disseminate existing data and information electronically. Second, a uniform system of medical terminology and language, as well as standard formats for the collection and reporting of laboratory and clinical data, must be developed. Third, new tools and techniques for processing data and information into knowledge bases for intelligent retrieval and decision-support services must be developed.

Strategies for future research

Specific recommendations to address these research needs are:

(1) Develop methods and tools for building knowledge bases, updating them as new discoveries are made.

(2) Develop techniques for identifying gaps and inconsistencies in knowledge bases and for maintaining their integrity.

(3) Develop new techniques for automatically extending knowledge bases by drawing inferences from raw data and by processing information available in textual form.

(4) Develop standards for medical terminology and language that can be used to define concepts and relations in the knowledge bases.

(5) Develop uniform standards for electronic transmission, dissemination, and sharing of imaging, acoustic, and instrument data, and graphic and textual information among various institutions and the Library.

(6) Gather laboratory and clinical studies for archiving of essential raw data by the Library.

Medical Decision Making

Rationale for research in the area

The term *medical decision making* encompasses some of the most significant research efforts of NLM. Broadly stated, medical decision making refers to any technology that can support a health-care provider or administrator in the process of evaluating options and making choices to optimize patient care in terms of outcome, risk, cost-efficiency, and patient satisfaction. The scope of the field is broad and includes:

- Deterministic strategies, such as clinical algorithms or ordered decision procedures, that guide decisions to one of a number of predefined outcomes when the data and constraints of a problem are well specified;
- Probabilistic strategies, such as Bayesian, decision analytic, and some expert system approaches, that combine detailed, formal representations of decision alternatives with quantitative estimates of uncertainties and likelihoods and values of outcomes along the decision paths to try to compute an optimal decision;
- Heuristic strategies, such as descriptive decision analysis and expert systems approaches, that attempt to emulate human expertise in constructing understandable symbolic lines of reasoning, taking into account quantitative, analytic, experiential, and uncertain information as appropriate, to arrive at an effective decision.

There is a growing need for the development of more effective decision-support technologies in medicine. The vastness of medical knowledge makes it impossible for a clinician to keep all relevant information in mind while making many diagnostic and therapeutic decisions. By understanding and formalizing the process of medical reasoning more completely, computer-based decision-support systems can help health-care providers solve daily problems. Expert systems can help locate and retrieve relevant information from a library or patient-care information system.

Decision-support systems can help to standardize care, optimize repetitive procedures, and provide a framework against which to assess nonroutine decisions. The resulting uniformity can aid in the collection of reproducible, reliable data on outcomes and efficacy, two areas difficult to control with present clinical research methods. Standardization can lead to efficient clinical protocols and methods for the evaluation of protocols.

Decision-support systems, once they are in routine use, have substantial potential for cost savings in automated surveillance systems that optimize scheduling and avoid unnecessary and duplicative procedures in hospital and outpatient treatment facilities.

There are many potential users of decision-support systems, including physicians, nurses, and other members of health-care teams during daily patient-care activities. Because of shorter hospitalizations and greater use of ambulatory care, patients and lay persons will become more responsible for their own care and thus will need decision-support technologies. Finally, public health officials and health-care managers will be important users. They will use the systems to manage health-care environments to assure the delivery of effective and cost-efficient care and optimal allocation of resources.

Recent accomplishments

Decision-support systems are receiving increasing attention and being more widely applied. Articles are beginning to appear in clinical specialty journals with some frequency, and a few experimental systems are being integrated into automated hospital or ancillary service information systems. The HELP system from Salt Lake City is one example; the TRIM system from University Hospital in Boston is another. Cost-effectiveness analysis has been added to NIH's Consensus Development Conferences, and decision-support approaches are being considered by medical groups that range from specialty societies such as the American College of Cardiology and the American College of Gastroenterology, to the World Health Organization.

Decision-support systems based on expert system technology have also made considerable progress, and operational systems are already being used in pulmonary function analysis and the analysis of electrophoretic patterns. Other clinical consultation expert systems are being used experimentally, INTERNIST-1 for diagnosis in general internal medicine and ONCOCIN for cancer chemotherapy protocol management; others have been used in earlier experiments, such as CASNET for diagnosing ophthalmological diseases.

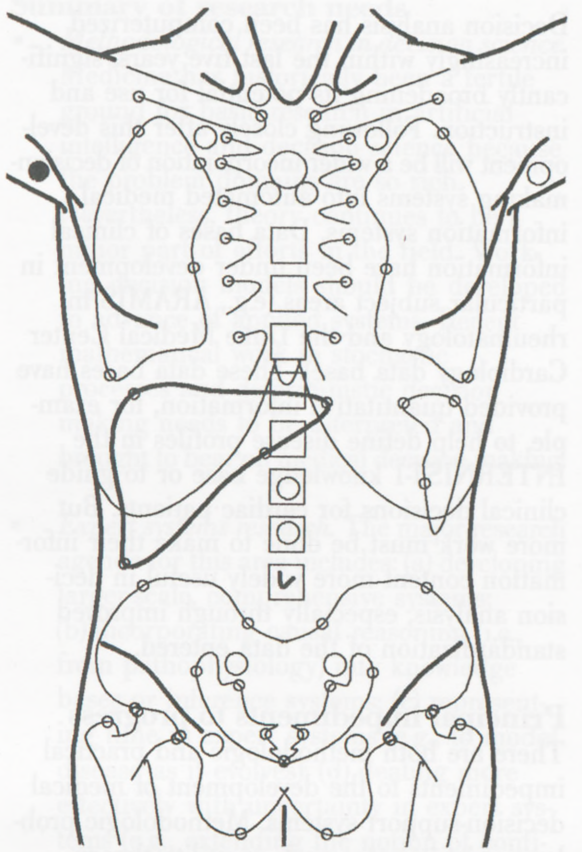

Graphic capabilities of modern computers are permitting new forms of data capture in clinical settings. The ONCOCIN system accepts graphical specification of areas of disease involvement by physicians caring for patients with cancer.

Other advances in decision-support systems have occurred in the areas of representing outcome utility for alternative patient therapy and management approaches, modeling of problems, and evaluating expert systems. Methods for assessing patient preferences have received significant attention in the last decade, as have the development of new metrics for decision- and cost-effectiveness analysis, non-Bayesian formalisms for incorporating the results of chance events, alternative stochastic representations for uncertain decision problems, and the incorporation of epidemiologic concepts into decision-support systems. Experimental expert systems are approaching medical decision problems not only from the viewpoint of consultants in constructing diagnostic or therapeutic decision rationales, but also from the viewpoint of critiquing hypotheses posed by users.

Decision analysis has been computerized increasingly within the last five years, significantly broadening its potential for use and instruction. Following closely after this development will be a wider incorporation of decision-making systems into automated medical information systems. Data bases of clinical information have been under development in particular subject areas (e.g., ARAMIS in rheumatology and the Duke Medical Center Cardiology data base). These data bases have provided quantitative information, for example, to help define disease profiles in the INTERNIST-1 knowledge base or to guide clinical decisions for cardiac patients. But more work must be done to make their information content more widely useful in decision analysis, especially through improved standardization of the data entered.

Principal impediments to progress

There are both methodologic and practical impediments to the development of medical decision-support systems. Methodologic problems exist in areas such as utility theory, where it has proved difficult to obtain assessments for the expected utility of differential life expectancies; problem formulation, where the sensitivity of decision models to different stochastic representations and types of statistical variability needs to be tested; and system evaluation, where the standards and methodologies to be used have not been rigorously defined or validated.

Practical problems are even more pressing.

- There is only a short history of the use of decision-support technology in day-to-day patient and health-care management. Vendors are reluctant to add unproven technologies to commercial systems, so development has been limited to single-institution research systems. Although there is an expressed interest in validating other scientists' work, the academic resources available for validating decision-support systems are small, and there are no established protocols for how to go about it. It may be assumed that the ultimate test of the validity of a decision rule or system is its transferability and utility from one clinical setting to another. But such transfer is very difficult at present because of differences in the computing and clinical environments.

- There is a lack of trained investigators and users of decision-support technology. The number of active basic researchers in medical decision making is no greater than 50, and the number of clinical expositors and clinical researchers, although larger, is inadequate to the task of educating students and trainees in over 125 medical schools. At present, no network of education and training exists for faculty in these areas, nor is there a mechanism by which nonexpert, peer role models can be fashioned outside the few leading centers.

- Decision-support and expert system development tools are still cumbersome. Workstations are just beginning to become affordable for the developer; user workstations are in the future.

- There are possible legal impediments to the licensing and use of decision-support systems in medicine. The Food and Drug Administration has claimed to be an appropriate regulatory agency for decision technology, but has proposed no formal mechanism of licensure. Although some preliminary essays have been written about the status of decision-support systems with respect to liability (e.g., malpractice, product liability), no precedents have been set that clarify their legal status. When systems migrate from institution to institution, questions of legal ownership and liability also arise. No work has been done on the questions of intellectual property inherent in knowledge bases, nor on the patentability of decision strategies.

- Physicians have been resistant to decision-support systems. Many have been unfamiliar with computer technology, disillusioned by the early extravagant claims of medical computer science, and uncomfortable with the possibility of losing control over decision making. As the true role of decision-support technologies becomes more widely known—consulting, not care giving—and as the current generation of trainees becomes more computer literate, this resistance should slowly abate.

- The widespread clinical usefulness of decision-support systems has not yet been demonstrated. These systems are still principally in the domain of research, and effective, integrated systems are still quite rare. Only when decision systems are exported from developers to other physicians can their clinical usefulness begin to be evaluated.

Summary of research needs

- *Methodological research in decision science.* Medicine has historically been a fertile ground for basic research in artificial intelligence and decision science because the problem domains are so rich. Nevertheless, theory continues to be a minor part of efforts in the field. Working decision models should be developed in advance of applied systems. Recent mathematical work in stochastic processes and deterministic decision making needs to be interpreted and brought to bear on medical decision making.

- *Expert systems research.* The major research agenda for this area includes: (a) developing larger-scale, comprehensive systems; (b) incorporating causal reasoning (i.e., from pathophysiology) into knowledge bases or inference systems; (c) representing time in expert systems (e.g., to model disease as it evolves); (d) dealing more effectively with uncertainty in expert systems (e.g., extending the notion of confidence factors for inferences); and (e) analyzing the tradeoffs between various forms of reasoning and the effects of such tradeoffs on the efficiency of strategies.

  As decision analysis and expert systems research appear to be converging, it is important to explore hybrid systems over the near term. Decision analytic models can be expanded to incorporate knowledge bases, causal reasoning, and temporal reasoning. Conversely, expert systems can expand their treatment of uncertainty.

- *Workstations.* Decision-support and expert consultation/critiquing system development tools need to be created inside powerful personal computational environments that have access to the diverse data and information bases important to decision support.

- *Integration of decision technology into medical information systems.* Approaches need to be developed to merge decision-support functions with the routine aspects of medical information and patient record systems, including, for example, the identification of problems in data surveillance and automated updating of knowledge bases and parameters in decision models. Prototypes of decision systems that monitor patient data and infer appropriate changes to practice should be developed and validated. Such automated decision-support systems will also need to be adapted to the user, and techniques will need to be developed that determine when the nuances of a problem place it outside the scope of the decision-support system.

- *Training.* Medical schools, nursing schools, and professional societies are wrestling with the need to incorporate computer-based systems into their educational activities. Decision-support systems are an essential part of such programs. Until entering students are uniformly able to use computers effectively, this issue must have a high priority.

Strategies for future research

NLM should continue to fund both expert systems and decision analysis research, especially when methodologic issues are under study. Hybrid systems should receive particular attention.

(1) Applications of these technologies to clinical systems, hospital information systems, and instructional systems should be highlighted. Nonintegrated efforts in medical decision support without a novel methodologic focus should be de-emphasized.

(2) The categorical NIH institutes should be recruited as funding sources for clinical research using decision-support and expert systems technologies. To an extent, the institutes are already involved in research in medical decision making. Examples include decision analytic clinical research, causal modeling, and expert consultation systems. Additional recruitment of specialty organizations and industry should be sought for applied research in decision science.

(3) Some developmental funding of workstations that are well integrated into clinical environments should be encouraged. In particular, demonstrations of the application of decision technologies to enhance the intellectual environment of medical practice should be proposed.

(4) Export and validation studies are high-priority items over the near term. Networks over which information that is common to transportable decision-support systems can be accessed are important supporting structures for this research.

(5) Networking of researchers in the informatics field is an important immediate priority for two reasons. First, it will significantly assist the development of this field. Consortia of AI (artificial intelligence) and other research groups exist; the decision scientists and the informatics educational researchers need to be similarly linked. Second, networking within informatics will demonstrate the feasibility of the approach in clinical disciplines.

Medical decision making provides a set of powerful tools that in many circumstances has been applied haphazardly to inappropriate problems. Future research in the area should demonstrate that the technologies are effective in the field, and that important changes in medical practice can result from the availability of online decision support.

Cognitive Issues in Medical Informatics

Rationale for research in the area

Diagnosis, treatment, and management of disease and health problems can be viewed as a series of problems to be solved and decisions to be made, involving both clinicians and patients. The information and knowledge on which advice and assistance are based is located both in the published literature and in the heads of the individuals involved. With the goal of improving health care, cognitive research in medical informatics aims to describe and analyze the processes of information search, storage, retrieval, and utilization and also the integration of knowledge that is based in the literature with that in the clinician's personal experience and understanding.

Skill in clinical judgment was long thought to be a fundamentally unanalyzable art, acquired through apprenticeship to skilled clinicians and personal experience. Recent advances in both experimental psychology and computer science have led investigators to take a new look at these issues. It has become possible to study processes of problem solving, judgment, and decision making more rigorously than before and to develop theories that yield greater insight into these complex mental processes. New knowledge has been gained about how both expert and novice clinicians process, evaluate, and respond to clinical information. On the practical side, the work aims to help experienced clinicians improve the quality of their decisions and to help novices acquire this expertise more efficiently by identifying more precisely what needs to be learned, in terms of both content and process.

In addition, as patients have become more involved in decisions about their care, it has become increasingly important to understand their strategies for seeking, evaluating, and integrating information.

Recent accomplishments

In the past 10 to 15 years, considerable progress has been made in our understanding of the diagnostic process in clinical medicine. It is now reasonably clear that diagnostic problems can be solved by a number of methods, including pattern recognition without conscious awareness of causal connections, hypothesis testing, and the use of detailed models of disease processes with hierarchically organized lines of reasoning.[35 36] All of these methods are concerned with the problem of how to provide focus and structure for a very ill-structured problem. Well-organized knowledge bases are needed to use these methods flexibly and efficiently. Experts know more than novices, but even more important, their knowledge is organized so as to be retrievable when needed and to permit shortcuts and efficiencies, by recognizing patterns quickly or pursuing only probable hypotheses and systematic lines of reasoning. Expert problem solvers can use deep understanding of the issues to combine steps instead of proceeding always in a uniform lock-step fashion.[37] Their knowledge base includes disease mechanisms, weighting schemes for evaluating evidence, therapies, rules governing the use of procedures, and higher-order rules and structures linking all these into understanding. Much of the knowledge about what to do next and how clinical findings should be weighted is not found explicitly in textbooks, but is informally transmitted during clinical training.

Clinical judgment must integrate information about the patient's particular history with more general propositions about the clinical problem and possible treatments. Errors in judgment can be caused by lack of information, erroneous information, difficulties in identifying relevant information, or errors in processing large amounts of information.[38] Since knowledge about diseases and therapies is expanding rapidly, and the growth of medical technology has made it possible to collect more information than ever before, we can expect these problems of information access, retrieval, integration, and processing to become more and more crucial. Much of the cognitive research in medical informatics is concerned with identifying the knowledge structures and reasoning strategies that characterize expertise, cataloging impediments to efficient information processing, and developing techniques to facilitate better use of information and health care.

Principal impediments to progress

Why would a health-care professional fail to use existing knowledge and information to arrive at the best possible decision? There are several reasons, all related to difficulties in obtaining the right information at the right time in the right form and quantity or to difficulties in integrating the information with prior knowledge and attitudes into a logical conclusion.

For example, clinical practitioners have severe problems in making diagnostic and therapeutic decisions because they must deal with a very broad range of problems and find the necessary information within fairly narrow time constraints. Yet the key points for a particular decision may be buried in a lengthy research report, or perhaps expressed in technical language outside the expertise of the decision maker. The effort required to locate and read the report, even assuming the reader will be able to interpret it correctly, may be prohibitively great.

Sometimes a clinician is overwhelmed by the sheer quantity of the published literature or the amount of clinical data that can be collected on a patient, if all the technical resources available were to be employed. How much information is enough? Edwards demonstrated that humans are often conservative decision makers, continuing to seek additional information long after they have enough to make a satisfactory decision.[39] This problem is further compounded by the conflicting findings of research studies. By using different populations and methods, separate investigations of the same research question may yield varying conclusions. When relevant studies are scattered in a number of different sources, not readily identifiable by title or key words, and reported in different formats with key points of comparison described incompletely or not at all, finding and critically comparing the important studies in an area quickly becomes overwhelming, particularly when the objective is a timely and practical decision rather than a scholarly dissertation.

Everyday clinical reasoning generally avoids using statistical principles to draw inferences. In particular, the interpretation of laboratory tests commonly disregards disease prevalence and its effects on test interpretation.[40] The probability of an outcome or event may be judged by how easily instances can be recalled to memory (how available they are) rather than by how frequently they actually occur.[41] Since statistical and epidemiological data are lacking for many clinical problems, the availability principle will then be used to generate estimates of risk (probabilities). Because it is easier to imagine a rare event than its likelihood, small probabilities will be overestimated, and the probability of a rare disease will be psychologically overweighted, compared to the probabilities of more common alternatives.

People feel more strongly about avoiding potential losses than achieving equivalent gains. Consequently, preferences can be markedly affected depending on whether a particular health-care outcome is judged as a gain or loss, relative to the current state of affairs. These judgments, in turn, can be affected by minor variations in the wording of the alternatives. Shifts in the focus of attention appear to exercise major effects on choice.[42]

Desirability of complex, multidimensional health states is often judged by reference to the most important attribute instead of a weighted combination of all features. In particular, the wish to minimize the chance of causing any harm may lead both patients and clinicians to forego higher chances of compensating gains.

Summary of research needs

Knowledge is useless for problem solving until it moves from the literature into the heads of providers and patients. More detailed understanding is needed about cognitive structures and representations in broader domains of medicine than have yet been studied. We know very little about the cognitive restructuring that occurs as learners move from being novices to experts in a domain, or about the effects of experience on information processing.

In the teaching of diagnostic inference, we are not yet able to give students much guidance as to how they should go about constructing a fruitful line of reasoning or set of hypotheses. How can we more efficiently avoid exhaustive searches for nonexistent diseases and yet not overlook rare conditions that are potentially treatable? Research is needed to understand more precisely the heuristics used in cognitive structuring of early hypotheses, problem formulations, or lines of reasoning. Systematic instruction in clinical reasoning and decision making is still in its infancy.[43,44]

Better strategies are needed for helping clinicians think about complex problems involving tradeoffs between risks and benefits. These problems will become more prevalent as the population ages and more citizens have multiple, long-term health problems. Focusing attention on a single problem, a common simplification strategy, is unlikely to be very effective in these instances. Decision analysis and related multivariate quantitative strategies have much potential, but their formal quantitative requirements may limit their appeal to broader clinical audience.[45] How can we do better in a practical way?

Research on the prevalence and seriousness of cognitive inefficiencies, biases, and errors in clinical inference has barely scratched the surface. Laboratory studies should be supplemented by more clinically oriented investigations to better relate principles established in laboratory experiments with clinical practice. The goal would be to identify the task conditions that evoke or minimize errors and biases, such as using the most readily remembered instances to estimate probability. By restructuring a problem or the way data are presented in medical records, articles, and textbooks, we may be able to help clinicians minimize errors in information processing and judgment. Another related question involves the factors that affect the focus of attention in clinical decision making. This has important effects on choices, as described in the previous section.

A deeper understanding is needed of attitudinal factors that affect the treatment preferences of both clinicians and patients. Both may be unreasonably optimistic or pessimistic. Some people feel they are inherently lucky, while others feel bad things always happen to them. These attitudes may affect the decisions people make for themselves and others. We must also recognize that in many clinical situations, there is as yet no clearly dominant therapeutic strategy, and that preferences in these situations ought to depend at least in part on the patient's attitude and life style preferences. Clinical decisions in these situations are not simply matters of choosing a technique; the medical information system of the future should clearly distinguish between questions of value and questions of fact and provide guidance about both wherever possible.

A better understanding is needed of how to integrate patients' and clinicians' values into the medical decision process. In several studies, assessment of values and preferences has been shown to be sensitive to the type of questions asked and the way they are asked, but there is no standardized way of doing these things, nor are we sure yet how pervasive and serious these problems are.

Research is also needed on how patients and their families think about their clinical problems and the advice they receive from professionals. This research is especially important as medicine and nursing increasingly turn their attention to the management of chronic diseases and the problems of an aging population. Crucial elements of a patient's attitude may not be expressed verbally, but may be communicated instead through body movements, postures, facial expressions, hesitations, silences, or behavior that seems inappropriate to the situation. Skilled clinicians traditionally give artful attention to nonverbal expressions and nonrational issues, combining this less formal information with scientific knowledge to arrive at appropriate decisions. For decision-support systems to take fuller account of information from these sources, research must address the following questions:

- To what degree can a clinician's intuitive understanding of nonverbal expression be formalized?
- What strategies are possible for promoting effective interaction between this intuitive understanding and formalized but necessarily simpler models of decisions contained in decision-support systems?
- Is it feasible to improve the quality of clinical decisions by creating human-machine dyads that capitalize on the unique strengths of each component?.

How can decision makers in clinical settings be assured of having examined and evaluated a sufficient quantity of valid, reliable, and relevant information in order to have a sound basis for their decisions? There is a need to develop knowledge bases better oriented to the information needs of practitioners, thus supplementing existing knowledge bases that are more suited to the needs and time perspective of the research community. These new knowledge bases should be practice oriented and user friendly, providing advice about what to do and how to interpret clinical evidence. They should be adaptable to different levels of experience and expertise, and should dispense advice and guidance just ahead of where the user is. They should also be geared to the information needs of all health professionals and patients.

Better methods should be developed to verify whether a user's questions have been correctly understood by computerized decision-support systems, and whether the user has correctly understood computerized presentation of information, both verbal and visual. The problem is analogous to that arising with human consultation and giving of advice.

Strategies for future research

NLM should stimulate and support research and training programs in the following broad areas:

(1) Useful problem representations and problem-solving strategies for more efficient use of medical information;

(2) Evaluations of the impact of cognitive and judgmental heuristics and biases, on health-care decision making and identification of task and user characteristics that affect these factors;

(3) Strategies to weigh risks and benefits, to assess values and attitudes of patients, and to integrate patients' and clinicians' values into the medical decision process;

(4) Information search and use by patients;

(5) Nonverbal communication between patient and clinician;

(6) Human-machine communication, including verification and evaluation research.

The research areas identified imply the emergence of a field that might be called medical cognitive/decision science. NLM should support training programs aimed at producing the next generation of scientists in this field. Extending and improving the quality of research in this domain is currently limited by the small number of experimental cognitive psychologists who understand enough about clinical problems and decisions to con-

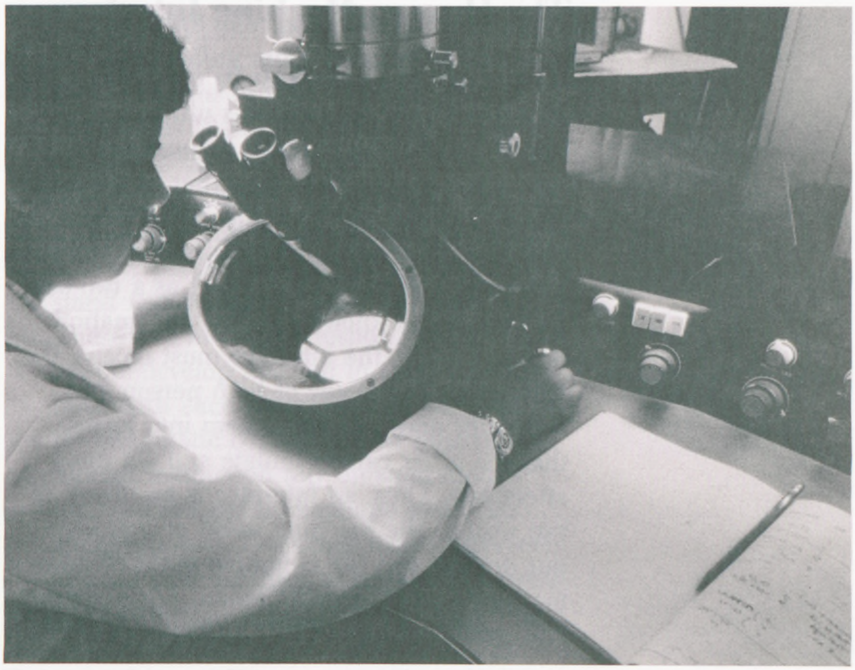

duct good research in these application settings, and the few physicians, nurses, and health professionals who understand enough about concepts, methods, and theories of cognition and decision making to conduct research that builds upon, tests, and expands theoretical knowledge. Consequently, interdisciplinary training programs should be especially encouraged. Support of training programs in this area would signify recognition that issues of knowledge representation, storage, retrieval, and utilization for problem solving and decision making are addressed not only by research on computer applications to these fields, but also by research on how human beings carry out these tasks and how human information processing and management can be improved. As such, extending research and training to this domain is consistent with NLM's long-range mission.

The Human-Machine Interface

Rationale for research in the area

Medical information systems ultimately involve interaction between a computer system and a biomedical researcher or medical decision maker. Between these two very different information processors is an interface with a difficult task: to act as a two-way transformer of knowledge and information needs. That is, the interface must receive the knowledge and needs of both person and computer and transform them into forms that the other can assimilate and act upon effectively. Thus, an interface is an object of two worlds: the machine world of input/output devices and the cognitive world of information/knowledge assimilation, organization, and presentation.

Accordingly, interface research in medical informatics addresses two categories of activity: issues of human-machine interaction (interactive devices, graphics, and the mechanics of computer use) and cognitive issues (user modelling, natural language interaction, and explanation). Since interfaces mediate between the medical decision maker and the computer, medical decision makers will not have access to the information they need, when they need it, and in a form they can use most efficiently, unless there are advances in these two aspects of interface technology, as well as concomitant studies of the effective use of other emerging technologies.

Recent accomplishments

Over the past 20 years, significant improvements have taken place in the interactions of humans with machines; they have progressed from batch, job control language interfaces using punched cards, to real-time, interactive interfaces using CRT (cathode-ray tube) graphics displays, a typewriter keyboard, a pointing device (''mouse'', light pen, touch screen), and command languages more natural to human use. These improvements have been possible only through the development of sophisticated software systems and hardware devices, many of which are still found only in research laboratories and some universities. Some of the important advances in HMI (human-machine interaction) include:

- The spread of personal computers (mostly into offices and homes) and the development of a pointer and icon-based style of interaction more suitable for computer-naive users. This style is eminently suited to the wide range of computer-supported tasks that involve a limited number of options: data base management, word processing, spread sheets, form-based interviewing with menus, pictorial displays of medical information and knowledge, etc. Many of these are relevant to medical informatics, and the style will probably also become incorporated into other knowledge-base management applications.

- The development of high-resolution bit-mapped displays, which allow at least an order of magnitude more information to be effectively displayed on a screen than was possible in the past using conventional CRT screens.

- Progress in 3-D (Three-Dimensional) reconstruction techniques so that data gathered by 3-D sensing devices such as CAT (Computer Assisted Tomography) scanners and 3-D NMR (Nuclear Magnetic Resonance) imagers can be interpreted, reconstructed, and stored as an accurate 3-D record of the imaged object.
- Some progress in the use of speech for human-machine interactions, both in the acceptance of spoken input to the machine (isolated word/phrase understanding) and in the generation of spoken output.[46]

In the cognitive aspects of interface technology, significant advances have also occurred:

- Natural language interfaces to data base systems now constitute a well-tried technology, and robust systems are now commercially available—e.g., Artificial Intelligence Corporation's INTELLECT,[47] Texas Instrument's NLMenu,[48] and various interfaces from Cognitive Systems and the Carnegie Group. These systems allow a user to retrieve information, even from very large data bases, by posing questions and commands in everyday English. Once installed, these systems are extremely simple to use, and allow even a first-time user, in a matter of minutes, to get relevant information that previously may have taken several days with a systems programmer as intermediary. Nevertheless, these systems have significant limitations that prevent their extension from natural language interfaces into more general interactive decision-support and problem-solving systems.
- Advances in systems' ability to generate well-formed, understandable natural language texts that explain an expert system's conclusions[49] [50] or define terms found in the data base domain.[51]
- Initial forays into the realm of modeling users, including attempts to represent users in terms of their previous knowledge, so as to present only as much information as they need and can understand in a comprehensible form. Attempts have also been made to model users in terms of their plans and goals, so as to understand and respond to their needs.[52] These user models correspond to those that people develop and use in their interactions with each other. In fact, many of the pragmatic conventions of human-human interaction rely on such modeling, and normal interactions would go awry without it. Thus, user modeling is seen as a valuable addition to human-computer interaction.

Principal impediments to progress

Difficulties with current interfaces to medical data and knowledge exist at several levels. The most obvious problem is that interfaces are not integrated into the work environment; the physical equipment that could provide an interface with needed medical data or knowledge is at present rarely located at the decision-maker's side at the point of decision making.

Another problem is that medical information systems are not integrated with each other. That is, knowledge is commonly distributed so widely through a host of mutually mute subsystems that an inordinate amount of effort is needed to jump back and forth among them, using awkward access methods to pick up the dispersed bits of information one needs. As noted earlier, where there is only a limited set of options, pointing devices provide a rapid way of stating one's needs.

With many options (e.g., in specifying symptoms, treatment methods, drugs, etc.), typing at a keyboard is still the only way. For many medical personnel, this is awkward, frustrating, and prone to error. Speech recognition is still not at a stage where spoken input—quick, direct, and natural—is possible.

Negative reactions to problems like lack of proximity and integration, as well as slow and awkward access methods, confirm the validity of what might be called the *inertial theory of information:* the use made of any piece of information is inversely proportional to the effort needed to get it. Even with vast improvements in the quality and range of medical information that can be provided through advances in knowledge representation and decision-support tools, this information will not be incorporated into daily decision making unless it is rapidly and easily accessible.

Difficulties with interfaces to medical data and medical knowledge also appear at the level of communication channels. These are problems of size; medical decision making in many areas requires simultaneous visual access to large amounts of information and the bodies, joints, and other objects of interest in medical decision making are complex and detailed and so require large amounts of information to be represented, manipulated, and displayed. Current hardware devices and software systems are not yet able to manage this information inexpensively, in real time.

An additional impediment to progress in interfaces for medical informatics systems is that much of the efficiency and effectiveness of human-human interaction (particularly, problem-solving or advisory interactions) come from two poorly understood aspects of natural language communication. First, some conventions of natural language use allow a large amount of information to be communicated implicitly. Second, a rich functionality allows both parties in the transaction to verify that they have understood one another; to recognize when they have not, and to effect a repair; to understand the boundaries of the current situation by positing hypothetical changes to it; and to ask for and receive appropriate explanations and justifications. Solutions to these problems depend upon successes in cognitive science, artificial intelligence (natural language processing), and linguistics.

Summary of research needs

Our recommendations for support of interface research focus on both HMI (human-machine interactions) and cognitive aspects of interaction. The improved integration of interfaces into the work environment involves the investment of both time and money into the design of specialty-related workstations and another large investment to bring them as close as possible to the medical decision maker and as easy to use as the nearest telephone.

Solutions to the problem of integrating subsystems with each other are appearing in specialized settings, as in clusters of consumer-oriented PCs or researcher-oriented workstations, but these have not yet found their way into systems used by practicing medical personnel. Improvements have taken place in intersystem communication to facilitate interchanges: these include terminal connections, file transfers, remote data base queries, and cooperative computing to share tasks. Special medically oriented solutions to such problems are appropriate areas for NLM-supported research.

The problem of providing simultaneous user access to large quantities of information (e.g., combining image data, graphics, and text) requires the development of larger high-resolution displays than are currently available. This is especially true in areas like radiology, where one must be able to simultaneously view many complex images, and in areas like medical expert systems, where one may want to simultaneously view multiple aspects of the system's decision-making activities. On the other hand, large displays have the same management problems as large collections of paper and pictures: how to move easily and flexibly among individual items in order to relate or compare them to others.[53] One needs to be able to focus on the information that is needed and to be able to select additional pages if necessary without shunting others out of visible memory. Research is accordingly needed in the area of intelligent display management, in line with advances in display technology.

Research is also needed on the representation, manipulation, and display of all kinds of 3-D objects of interest to medical decision makers. For this purpose, decision makers must be able to visualize and interact in real time with representations of objects, including the ability to request 3-D information in a natural manner (e.g., the anatomical correlates of penetrating wounds or complete views of 3-D objects),[54] the ability to model nonrigid bodies in a semantically significant manner (e.g., modeling the distribution of organs within the torso), and the ability to modify displays of 3-D objects in real time (e.g., for surgical planning).[55]

A great deal of research and development will be needed to equip machines with the capacity to deal competently with natural language, both when input by human users and output to them, for use when simpler graphical interaction techniques are inadequate. Expert systems for diagnosis or patient-management advice should be able not only to produce the correct recommendations but also to explain them. An ideal system would recognize and correct user misconceptions as they were revealed in the course of human-machine interactions, much as a good human tutor or consultant does. Computer programs that can let users verify their understanding of the system's advice, and do so in natural language, would be a significant step forward, as would techniques for giving users more leeway in responding to the machine's requests for information. A valuable model for understanding some of the requirements of question answering in medicine has been described.[56]

Research is needed on speech understanding and speech production, both of which will have important uses in medical informatics. Several commercial systems now can recognize and respond to spoken commands and can produce recognizable spoken letters, numbers, and words or short phrases. But these systems have not yet achieved vocabularies of adequate size nor easy use for continuous speech recognition or production in natural language patterns. Research in this area will be difficult and lengthy, but will lead to more flexible and natural interfaces to medical information systems.

Finally, research is needed in the area of user modeling: studies should examine the kinds of models needed for various aspects of interface design, how those models can be set up a priori or acquired interactively, and using those models to improve users' abilities to get the information they need and understand the information that they get.

Strategies for future research

We recommend:

(1) The design of specialty-related workstations geared to the information needs of particular medical specialties;

(2) Work on more effective integration of workstations with other computing and information resources;

(3) Development of even larger high-resolution displays for medical information and of techniques for intelligent display management;

(4) Development of real-time methods for representing, manipulating, and displaying diverse 3-D objects;

(5) Development of natural language capabilities to handle those aspects of interacting with expert systems and decision-support systems that cannot be carried out through simple icon and pointer-based interactions;

(6) Studies on the use of user models in improving the quality of interactive systems; and

(7) Development of speech understanding systems and voice-based information delivery systems.

Information Storage and Retrieval

Rationale for research in the area

The explosion of biomedical knowledge has compromised the ability of individuals to keep up with new developments and discoveries in health science research and clinical practice. To describe the growth in concrete terms, the first volume of *Index Medicus*, published in 1880, contained about 17,000 citations from 700 periodicals, while in 1986, a volume for a single month has about 21,000 citations from more than 3,000 journals and publications.[61] The current NLM collection comprises approximately 2.15 x 1012 characters and the volume of the collection would require over 4,000 of today's magnetic computer disk storage units to hold it—and this does not include any of the clinical and research laboratory data generated each year.

The problem for medical informatics is how to organize and structure knowledge, store and retrieve it as needed, and apply it in an efficient, efficacious, and cost-effective fashion.[57] Historically, the biomedical knowledge base has been organized in printed form at five levels, each synthesizing and abstracting the ones below it. These include raw research or clinical data, journal articles or research reports, review articles, specialized monographs, and textbooks. This traditional means of organizing knowledge is becoming inadequate for several reasons. First, the growing volume of information threatens to swamp the established library mechanisms. Second, it is increasingly difficult for investigators to locate and retrieve relevant information from other studies to integrate it with their own work, and for reviews, monographs, and textbooks to adequately capture and synthesize all the knowledge at the underlying levels. And finally, the rapidly changing knowledge base, coupled with the inherent time, effort, and cost involved in updating information in printed media, makes it increasingly difficult to keep the recorded knowledge base current. New approaches are needed to the management of this information, along with more *intelligent retrieval systems that will facilitate* access to these large data bases and draw from them only the information needed for a particular application.[58]

The library must move from being primarily a warehouse of passive knowledge sources to becoming an *active* collection of knowledge sources—one in which the knowledge is available for use by the inferential processes of computer programs. We must develop new hardware and software tools for the storage and dynamic management of the full range of medical information. This would facilitate the synthesis of knowledge when and where it was needed and would reduce the reliance on prepackaged syntheses. It would increase the richness, completeness, and relevance of the information that could be applied to the particular problem under investigation. Research on these goals is squarely in line with NLM's mission in the management of biomedical knowledge.

Recent accomplishments

NLM has been the leader in developing and facilitating operational access to biomedical bibliographic data bases. MEDLINE has been in use since 1971 and provides online query support from NLM machines and several commercial data base services for literature references from over 3,000 biomedical journals and publications. In addition to offering free text searches, NLM pioneered the development of a standard set of index terms MeSH (Medical Subject Headings) to facilitate the organization, indexing, and retrieval of the broad knowledge base in medical literature. The user interfaces for retrieval systems are improving, in that they no longer require the user to know the details of the data base organization and allow queries to be expressed in forms natural to the user rather than to the data base system. As knowledge bases grow, more intelligent interfaces will become essential to help the user locate information from diverse sources and subject areas relevant to research needs.

A significant amount of work has also been done on the problems of automating the storage and retrieval of patient medical records, part of the raw data of medicine. This work was largely directed toward solving the needs of individual institutions or health-care environments. Such systems offer investigators large data bases of patient information on which to conduct both prospective and retrospective studies. But because they represent mostly local collections of information, such studies will be limited until the data bases are linked nationally, providing the foundation, for example, for broader scope decision support systems.[59]

A few efforts have attempted to link decision-support systems to the literature base. In nursing, the COMMES system uses knowledge from selected undergraduate nursing texts and articles to support decisions about the usual nursing care of patients with common problems. The system differs from other decision-support systems in not operating from an extensive clinical data base, but rather from limited patient data entered on an ad hoc basis in an interactive consultation mode. The system relates the data to the literature base and offers suggestions about care and references for further study.[60,61] Other decision-support systems operate on a specially constructed clinical data base and cannot provide the user with ties to the broad literature.[62]

The literature base itself is changing with developments in electronic publishing. As more of the literature is made available in electronic form, the roles of the library in providing it and the user in accessing it are changing.[63] Libraries are evolving new roles to deal with these new developments, which will go beyond the passive warehousing of printed information to a new dynamic role of managing and facilitating information access.[64 65]

Much research is also underway to harness new data-recording hardware technologies, such as optical disks, for large-volume data storage on computers. These systems should become more practical in the next few years, and will offer 100 to 1,000 times more storage and support multiple media, including text, graphics, gray-level images, and audio.

Principal impediments to progress

A key impediment to progress in the new storage and retrieval methodologies is the still incomplete development of effective computational representations of medical information in ways that capture the rich semantic interrelationships. This topic is dealt with in more detail in the earlier section on knowledge representation. These representations provide the basis for the dynamic acquisition and synthesis of new knowledge and for the intelligent retrieval of relevant information in the context of particular health-care and research needs. The key will be a canonical language system—describing and relating knowledge about a myriad of topics—from patients, to pathologies, to biochemistry, to diagnostic and therapeutic strategies, and to cost-recovery mechanisms. Close behind is the enormous problem of converting the current vast store of biomedical knowledge, mostly in printed media, to an electronically accessible form.

Since much of the raw clinical and laboratory data of biomedical research will be distributed among computer systems that support individual groups, methods must be developed to describe the contents of these data bases and capture all of the contextual information that will make these data useful to related research and decision-support activities. These methods must make it possible to translate and communicate information among systems. They will require research on very large, distributed data base systems that will flexibly facilitate the encoding of complex semantic relationships and frequent updating. They must also appropriately account for protection of confidentiality and scientific attribution. And at the lowest level, protocols for accessing large data files must be developed that accommodate the diverse physical and logical structures of the data files stored in different machine environments.

Communication network connections, suitable for data transfer and sharing, currently exist among only very small parts of the biomedical research and health-care communities. The lack of such links makes it difficult to collaborate scientifically or aggregate or share electronic information. High-performance network connections must be commonplace throughout the broad biomedical community.

In the near term, 'smarter' retrieval systems must be developed with interfaces that better accommodate the background, context, and needs of the biomedical user.

Finally, significant progress is needed in the development of mass storage hardware technology and the software to manage it effectively for knowledge-based storage and retrieval systems. The technology routinely available today requires thousands of units to handle the large data files involved; it is clearly inadequate. Although it is unlikely that all of these data will be stored centrally, even distributed systems will require storage capabilities that are several orders of magnitude greater than those available at present.

Summary of research needs

This summary of needed research is organized into two areas, software and hardware. Although these areas are listed separately, they must be pursued in parallel to address the research goals of storage and retrieval technology.

Software Development:

- Develop representation and description methodologies for the large volumes of biomedical data and knowledge so that they can be stored in computational structures that facilitate the addition of new information, the resynthesis of existing knowledge with new information, and the retrieval of application-relevant information and knowledge.
- Define a data description language and extend communication protocols to facilitate communication, translation, and sharing of distributed data bases.
- Develop means to convert current data bases into forms usable in new storage media and retrieval system environments. In order to make use of information already available, there needs to be a way to bridge between the old systems to the new. The existing data cannot simply be retyped in.

Hardware Development:

- Develop higher capacity mass-storage devices that are able to manage the large volumes of raw data, including text, graphics, multidimensional images, and audio, and to facilitate the indexing, search, and retrieval of relevant information under software control.
- Adapt more powerful computers as needed to handle the large volume data bases and to run the intelligent search and information management software systems. This will probably require research in distributed and parallel computation to meet the needs of these applications.

Strategies for future research

Recommendations for reaching the research goals in this area include:

- Fund individual research efforts focusing on the underlying basic issues of biomedical knowledge representation, description, and storage. At appropriate points, efforts should be made to convert portions of existing information bases to new representation schemes.
- Fund prototype projects to investigate the integration of diverse information data base systems for specific application areas, such as clinical decision-support research. These should include components for data base management, use, and extension among diverse and geographically distributed user communities.

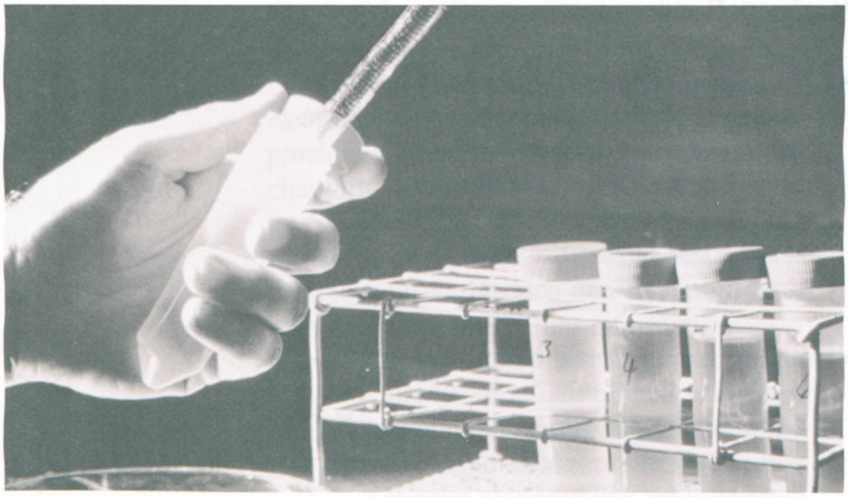

- Fund centers to investigate the development and integration of new hardware and software technologies important to large-volume biomedical information storage and management. This should include documenting the long-term hardware and software requirements of biomedical information storage and retrieval to stimulate industry and allied research activities.

- Establish electronic network communications through the biomedical research and health-care communities so that information dissemination and scientific collaborations can take place rapidly and conveniently.

- As the new methodologies permit, fund efforts to transfer the large volume of existing information bases into the new elec tronic storage and retrieval environments.

Technology Transfer and Dissemination

Rationale for research in the area

The need for transfer and dissemination touches all facets of medical informatics research and applications. In the development of medical information systems, we need the ability to transfer and share knowledge bases and problem-solving methodologies among research groups, so that they can compare alternative approaches and cooperate on the development of knowledge bases per se. It is equally important to disseminate integrated application systems within the biomedical community for further development, evaluation, and use. Finally, we need to transfer knowledge about the science of medical informatics itself to train new generations of researchers and users of these systems.

Attempts to transfer prototype systems to new operational environments are often frustrated by current incompatibilities among computer systems. A program that loads and runs on one computer cannot be loaded or run on another unless it is compatible with the new hardware and operating system. These intersystem transfer problems are further confounded by differences in the communication protocols in use on various vendor networks that link workstations among themselves and to other computing resources.

Just as de facto standards for compatibility seem to be developing—as reflected in the number of different manufacturers that claim IBM PC or UNIX compatibility for their machines—we need to begin to evolve standards for applications and equipment in medical informatics. One means of moving in this direction is to gain a consensus among funding agencies and researchers for generic standards for the types of operations and interfaces expected of machines and operating system environments, programs, and communication links among sites. Such standards need not impose constraints on the systems being developed by individual researchers, other than to establish protocols by which they can communicate their work in working form to other sites and machines. A simple example is the use of the ASCII (American Standard Code for Information Interchange) representation of textual information that enables several authors to collaborate on a common document (such as this one) using different word processing packages on different computer systems, even though the internal representation of text may be very different on each of the systems. Another example is the ability of modems to connect with each other, even though manufactured by different vendors. The standards define a common interface for exchange of data through which each individual machine "knows" how to communicate.

At a higher level, biomedical knowledge bases will accumulate over time and methods will be needed by which knowledge bases at various development sites can be interrelated and assimilated into larger sets of knowledge and shared among larger groups. This allows more development of medical knowledge bases to be distributed better and, it is hoped, more efficiently. It also allows researchers in several sites to study the efficacy of different decision-support and problem-solving approaches using the same putative knowledge.

This process actually happens in stages. When research on a particular application system is young and exploratory, it is quite appropriate that a prototype not be disseminated for wide use. As the work matures, it becomes desirable to transfer problem-solving strategies and knowledge bases to sites other than the primary development site in order to gain insights and critiques from other researchers to generalize the results. As the application matures further, transfer is necessary for field evaluation and validation and, finally, for wide dissemination of the work to the larger biomedical community and possibly the commercial sector. Routine methodologies for these kinds of disseminations must be developed to encourage essential cross fertilization of work, ideas, and training.

Recent accomplishments

Dissemination efforts to date have successfully experimented with three alternative approaches: shared access to systems on central machines, duplication of the hardware and software environments needed to run systems at other sites, and implementation of systems in languages that facilitate portability among different machines. There are numerous examples of nationally shared central hosts, such as MEDLINE, the PROPHET system for pharmacokinetic studies, the SUMEX-AIM system for artificial intelligence in medicine, the Chemical Information Service for chemical structure analysis, the GenBank system for molecular biology data bases, and the BIONET resource for symbolic computation in molecular biology. These systems are typically operated on mainframe systems at a major development or contract operations site, are connected to one or more of the national communication networks (e.g., ARPANET, TYMNET, CSNET, BITNET, TELENET, or UNINET), and offer time-shared computer access to the facility services. Many more such central systems provide more local support to university medical centers or regional communities.

Dissemination efforts on compatible hardware have been difficult and expensive until recently, especially for systems that require large mainframes. Still, long-standing examples of this approach exist, such as the SECS system for chemical synthesis and the PROMIS hospital information system. With the recent rapid fall in hardware costs, systems that run on commonly available equipment can be exported to remote groups. Other systems that are able to run on even less expensive personal workstations are even more easily exportable, at least within local communities. For example, the INTERNIST-1 knowledge base for internal medicine, coupled with a simple search program, is being moved to an IBM PC for easy export to other groups having the commonly available PC. Similarly, the HELP system for decision support and the ONCOCIN system for cancer chemotherapy protocol management are being developed on workstations.

Finally, some systems are written in languages that can be implementated on several hardware systems. For example, the DENDRAL chemical structure elucidation programs were rewritten from Lisp into BCPL to achieve portability to many hardware systems. Similarly, the Rutgers EXPERT system was written in FORTRAN for portability and others are written in PASCAL or C for the same reason.

Principal impediments to progress
There are a number of principal impediments to transfer and dissemination of the products of medical informatics research and development:

- Lack of standardization in programming languages and computing environments now in use by developers. Current prototypes are developed in a wide range of languages, on a wide variety of machines, and with generally idiosyncratic interfaces, greatly inhibiting dissemination. Of course, the promulgation of standards is in many ways an anathema to independent researchers, and one must be careful that standards are not created too early, freezing development and constraining necessary creativity and innovation.

- Lack of widely available networking among application sites. The ability to access expertise, knowledge bases, and machines across multiple sites for system development, applications tailoring, and consultation is needed to greatly hasten dissemination of programs.

- Lack of standardized interfaces with preexisting clinical records systems. Most sophisticated applications (e.g., HELP, ONCOCIN) are highly integrated with a hospital or clinic-based records system. Exporting these to new sites requires similar integration at the new site, where the existing records system is usually incompatible. Wide dissemination of most systems for decision support will depend on development of a standardized method for interfacing with hospital and clinic information systems.

- Lack of evaluation standards for testing systems prior to dissemination. Many systems currently under development are diagnostic systems whose goal is either to improve on expert clinicians' abilities or to replicate the ability of expert clinicians to make diagnoses. Since many diagnostic areas have no ultimate criteria of diagnostic correctness, the current standard of validity is expert opinion—and experts often disagree. As long as agreement with experts is the standard of evaluation, there will be room for improved methodologies to assess the impact of such systems.

- A lack of research and development manpower and training for researchers and end-product users in medical informatics. There are simply too few trained researchers in the field, and most now are devoted to working on system development. Creative and trained personnel are greatly needed to attack problems in the transfer of this technology. At the same time, sufficiently trained individuals are needed to create a target pool of applications managers—persons who can organize and support clinician-users at institutions that are not primarily involved in systems research.

- Lack of glamor and incentives in transfer and dissemination problems, compared to working on initial system design and development. The best way to attract researchers to these problems may be to target program funds to these activities and reward publications in this area.

Summary of research needs

The research needs to facilitate transfer and dissemination of medical information systems derive from the barriers identified in the prior section. They can be classified in two groups, theoretical and technological infrastructure.

Significant progress toward the goals of designing effective knowledge-representation methodologies must be made (see the section on knowledge representation). This will lead to ways of describing the contents of knowledge bases to facilitate their incremental development at different sites. It will also facilitate interfacing and evaluating them with decision-support and other problem-solving systems for which they were not initially intended.

Better evaluation frameworks must be developed to measure the impact of medical information systems of various sorts on medical care.

A consensus is needed on appropriate standards for implementation languages, user interfaces, network communications protocols, data base descriptions, and other elements, consistent with standards of other agencies and vendor groups, to facilitate system transfer.

An infrastructure of computer networking among health-care institutions needs to be developed to facilitate scientific collaboration, data and knowledge base development, and system sharing. This is a major resource commitment, but a necessary one. This is not to suggest the independent development and operation of a network by NLM, but rather to actively encourage cooperation with other agencies in connecting the biomedical community to selected Federal and commercial network systems that can provide needed linkages and support prototype developments to demonstrate the capabilities, resource requirements, and interface protocols needed for dissemination and use of medical informatics applications.

Strategies for future research
To promote transfer and dissemination of work in medical informatics, the following high-priority areas are identified and should be supported:

(1) Research efforts on knowledge representation must be aggressively supported, as should efforts to apply these results to promote distributed knowledge-base development and use.

(2) Research efforts should be supported to define appropriate evaluation methodologies for various types of medical information systems, such as diagnostic and decision-support systems and to apply the resulting tools to assess the impact of support systems on delivery of health care.

(3) Steering groups should be established to develop appropriate standards for medical informatics applications.

(4) The broad-based connection of medical informatics research groups to local area and national network systems should be supported.

(5) NLM's ongoing training program should be expanded to assure adequately trained medical personnel for the development and dissemination of medical information systems.

Supporting Technologies and Enabling Activities

Rationale for research in the area

A powerful, highly integrated, and widely available armamentarium of technological and human resources is essential to achieve the 20-year vision of medical information science applications described in the previous section, A Vision of the Future. New hardware and software computing resources are needed to collect, store, manipulate, and display the diverse kinds of information used in biomedicine and to develop and run the intelligent programs that will assist with patient care, research, library retrieval, clinical and laboratory instrument management, scientific communications, administration and financial management, and more. Extensive communication systems are needed to uniformly link health-care personnel to each other and to the diverse equipment and information resources that are becoming essential parts of their professional lives. Additional well-trained human resources are needed to develop, adapt, and employ these technologies for biomedical use.

For many years, medical informatics research has taken advantage of the rapid evolution in computing and communications technologies[66 67 68 69]—making available to all biomedical scientists more powerful, compact, and inexpensive hardware and guiding the development of more intelligent and effective biomedical software and communications systems. There are many diverse elements in the technology base that power medical information systems, ranging from semiconductor materials science, to computer architectures, to natural language understanding and knowledge-based systems.

Industry projections forecast continued reductions in the cost and increases in performance of microelectronics.[70 71 72 73] During the foreseeable future, for example, for constant performance, one can expect subsystem costs to decrease approximately 40 percent per year for memory, 35 percent per year for central processing, 25 percent per year for disk storage, and 10 percent per year for communications, while power and packaging costs remain constant. These hardware improvements, however, are only part of the picture. Software development is equally important, and a great deal of work has been done on academic research systems and, more recently, on commercially supported systems to assist with some aspects of the construction of large information management and problem-solving systems. The development of more and more capable hardware and software tools makes it clear that even more powerful systems will become available for medical research and practice each year.

As in the past, most of the basic computing technology needed for medical information systems will continue to develop outside of biomedicine, through support by the Department of Defense or industry. In some areas, however, external priorities for developing required technologies will not coincide with the needs of biomedical applications. For example, the military is funding research and development for problems such as battlefield management and aircraft automation, and industry is targeting product development for identified large marketplaces (e.g., IBM PC equivalents). In other technology areas, progress simply will be too slow, and in still others, help cannot be expected from outside sources at all, because the technology is specialized to biomedical application needs. This is particularly true in such areas as natural language processing, large knowledge-base representation and management, and medical problem-solving methodologies.

Special technological development will also be needed in the processes of importing, adapting, and making available technologies for specific biomedical purposes, such as nuclear magnetic resonance image enhancement and interpretation, physician/computer interfaces, clinical decision-support methodologies, knowledge-based medical library query systems, and the pervasive electronic communication links needed between computing systems and hospital, clinic, laboratory, and library information sources.

Finally, biomedical resources must be allocated to alleviate the substantial shortage of personnel—physicians, nurses, medical staff, librarians, and computer science and engineering professionals—who must be trained in the interdisciplinary technologies needed so they can lead and carry out the development of new medical information system applications.

Consequently, an effective program is required to provide essential supporting technologies to meet the needs of medical information system research and dissemination. NLM must play an active leadership and funding role in this effort. The goals of this effort are to help develop, adapt, and harness important technologies for use in research on medical information systems and to assure that the computing, communication, and manpower resources required for biomedical applications are available as needed. Given its limited budget, the focus of NLM's support must be clearly defined within the near- and long-term requirements of its research goals and balanced with the scope of support in place from other agencies and industry. In addition to enabling basic medical information system research, such support will ultimately stimulate commercial investment in further development.

Recent accomplishments

Much has been accomplished in adapting computing technology for biomedical research and health care, and NLM has played a leading role in this work. Computers help with tasks such as laboratory instrument control; data collection, storage, display, and interpretation; decision support and therapy planning; record keeping and data bases; financial management; mathematical modeling and statistical analysis; bibliographic search; text preparation and scientific publications; and collaborative communications and information sharing.[66 74 75 76] This work has also demonstrated the power of electronic communication tools to facilitate scientific collaboration and interpersonal communications, information and software sharing, and intermachine communications and centralized resource sharing. And the training programs in medical informatics supported by NLM have been the only manpower training programs in this area.

Until the middle 1970's, most medical information systems were built on mainframe computers, such as MEDLINE, PROPHET, and SUMEX-AIM. Following the introduction of laboratory minicomputers and the earliest personal computer workstations,[70] an increasing number of medical applications have moved onto such dedicated systems. Today, powerful individual user workstations are routinely and affordably available and dramatically affect all fields of scientific endeavor. They offer the medical practitioner and researcher an interactive computing capacity with graphics available only on expensive mainframes a few years ago. For example, systems like HELP and INTERNIST-1 run on IBM-PC-like workstations and the ONCOCIN cancer chemotherapy advisor now runs on a stand-alone Lisp workstation. At the same time, mainframe computers (with both serial and parallel architectures[71] and peripheral devices[77]) have become more powerful for large-scale numeric computations, data base storage and retrieval, and general shared use. Together, distributed workstation and central mainframe systems, coupled with high-performance local-area and national communication networks, offer an effective and essential set of computing resources for biomedical research and clinical practice.

Principal impediments to progress

There are a number of underlying technological impediments to making progress toward the long-term goals of medical informatics. First, scientific workstations lack needed capabilities and are still too bulky, noisy, and expensive. For example, we envision the workstation of the future to be a system that is the size of a clipboard or notebook, with 10 or more times the central processor speed and memory size of today's workstations, costing only about $1,000, and offering an effective and comfortable human interface. The least expensive workstations of today are very limited in processing power and memory. The most powerful workstations still cost from $50,000 to over $100,000, and even these run the most complex of today's programs with intolerable delays. With approaching limits of serial machines, even in high-performance, very-large-scale integrated circuit implementations, it is likely that parallel computation will be needed to achieve the computing capacities required for future systems—and very little is known about how to exploit large-scale concurrency in computer systems. Most of the current workstations have display/keyboard systems that are far from compact and portable. The operating system and software tools available still require expert computer knowledge on the part of the user and the modalities of interaction—speech, touch, vision, and multi-dimensional graphics displays—are very limited.

Second, only small segments of the biomedical research community have access to the integrated computing and network communications services that are essential to future medical information systems. A growing number of groups have modem access to networks like TYMNET, UNINET, CSNET, or BITNET that allow some terminal connections and simple electronic mail services. But these do not have high enough bandwidths or have the protocols in place needed for large file or data base transfer or for more sophisticated machine-to-machine interactions characteristic of a distributed computing environment. Also, few medical research and health-care groups are hosts, or providers of information services, on such networks. More capable research networks such as ARPANET or local area Ethernets also support file transfers, distributed file access, remote procedure calls, and other information exchange services that significantly facilitate research work and communication among computing and information resources on similar networks. But because they generally support computer science research activities, such networks lack connections with many of the information sources essential to medical information systems. Almost nowhere are hospital information, clinical laboratory, diagnostic and therapy specialty groups, library, research, and other biomedical resources conveniently and uniformly accessible by electronic means—locally, much less nationally. As these additional information resources become accessible by network, much higher data transfer and interaction speeds will be needed than are supported on existing high-performance networks.

Third, we have limited tools for the development, verification, and maintenance of large computer programs and knowledge bases such as are involved in medical information and decision-support systems. Today, these systems are developed with time-consuming hand crafting at all stages, from system conceptualization to design, implementation, debugging, and verification. We have very limited methodologies to verify that the implementation of a system conforms to the design intent or to the standards of expert medical practice. And changes to these systems become progressively more difficult as they are extended. The language systems used for today's software are often specialized to particular hardware and operating system environments, so that portability and sharing of programs and information bases is difficult and again, manpower intensive.

Finally, there is severe shortage of well-qualified personnel for the necessary development and dissemination of medical information systems. This shortage applies across all disciplines in which computer technology is being applied to the intelligent management of information. But the shortage is especially a problem in medicine because the needed cross-disciplinary training is so specialized and difficult. Physicians, nurses, other health-care professionals, and medical librarians must be trained in medical informatics. At the same time, computer science and engineering professionals must be trained to understand the unique requirements of biomedical information system applications. It is often easier for young engineering graduates to become established in a field directly allied with their engineering training (e.g., system design in military or industrial applications) than it is for them to become knowledgeable in medicine and address the difficult scientific, social, legal, and funding problems involved in biomedical applications research.

Summary of research needs

In general, the development of the technology base for medical information systems should be driven by the requirements for such systems as derived from experimental research and development and prototype evaluation. The principal areas requiring continued research include:

- *Hardware.* The development of more effective computer workstations for the broad range of medical information system applications must be encouraged. Future workstations must be faster, have larger memory, be more compact, have improved modalities for human interaction and information display, and be much less expensive than current machines. Hardware designs should anticipate that these systems will be ubiquitous and coupled to each other and to shared resources over many communications paths, often cooperating with each other through parallel processing. Cost-effective and high-performance systems are required for the storage, retrieval, and display of large amounts of diverse information, including text, speech, graphics, instrument data, multidimensional images, etc.

- *Software.* The development of improved methodologies for designing, implementing, and verifying large software systems is essential. For systems of the scale relevant to biomedicine, formal methods of program generation and correctness verification will likely prove intractable in the near term. Rather, the program development process is itself a knowledge-based activity, and expert systems to help manage large software systems will be needed. Improved higher-level languages for expressing and manipulating knowledge and solving problems must be developed. And finally, improved system software is needed that provides a more convenient and intelligent user interface to computing environment resources, remote graphics access to medical information systems, and efficient distributed and concurrent processing among workstation systems.

- *Communications.* Widely disseminated medical information systems will require high-bandwidth communications to allow access to the computational, data, and information resources needed for health care and research. These must include all of the broad range of resources involved, including hospital, clinic, laboratory, bibliographic, managerial, and research services. The software tools supporting these communication links must facilitate effective interactions among health-care providers and researchers and the computing environments that enable their work. Current communications networks, like the ARPANET, connect many computer science research groups together and could link biomedical resources as well, but only a limited effort has been made to do so. Such national communications systems, coupled with higher-performance local networks, will provide the essential 'glue' that binds together the various portions of the medical information science and biomedical research community.

- *Training.* The development of medical information systems, from the underlying basic technology to the actual application systems, and finally to their dissemination into health-care practice and research, will require substantial increases in the personnel trained to work in this highly specialized area. Increases are essential across the board in the numbers of physicians, nurses, medical librarians, and other biomedical, computer science, and engineering professionals with appropriate interdisciplinary training to meet the needs of academic research and teaching and eventually, to support industrial development of commercial medical information systems.

Strategies for future research

NLM's mission in medical informatics makes it a logical leader in the development of supporting computing technologies, but in very real terms, its budget is extremely modest in comparison with those of most defense agencies. Thus, NLM should be active in supporting the basic research that will fuel the development and dissemination of medical information systems over the next 20 years, but it must adopt a careful strategy to maximize the relevant return for its limited investment. The recommended approach includes:

(1) Adopt a "top down" approach, funding prototype efforts to develop, evaluate, and understand the detailed requirements for medical information systems and make these known to academic and industrial groups in biomedicine, computer science, and engineering research and development. These efforts should focus on the problems surrounding the design, implementation, dissemination, interface and integration, and maintenance of large medical information systems.

(2) Fund centers of excellence to monitor hardware, software, system, and communications developments, deciding opportunistically where strategic investments in basic developments would most benefit NLM programs. These centers should act as technology bridges to work being supported by other agencies, so that new and relevant technology can be quickly imported, adapted, and integrated for biomedical use where possible.

(3) Within available resources, fund basic computer science and engineering research relating to the design, implementation, verification, dissemination, and maintenance of large medical information systems. Candidate software research areas include software engineering tools; distributed concurrent system capabilities; and basic artificial intelligence research on knowledge acquisition, representation, problem solving, explanation, and natural language. Hardware areas include multimodality user interface devices, compact workstations, 3-D displays, and peripheral equipment for the storage and display of diverse large-volume information in forms such as text, speech, image data, instrument data, and graphics.

(4) Seek to establish standards, where beneficial, for medical information systems, in such areas as human interfaces, programming languages, data and knowledge-representation descriptions, and communication and information-transfer protocols to achieve better portability, communication, and synergy between research efforts.

(5) Ensure that supported medical information science research groups have access to state-of-the-art computing resources and to the software and communication tools necessary for effective work and collaboration. This should include establishment and support of communication and information-sharing networks among medical informatics researchers; other members of the biomedical research, clinical, and computer technology communities; and the diverse information resources needed for medical information systems.

(6) Continue active support of student training programs to increase the number of personnel—both medical and computer professional—available to develop and disseminate medical information systems.

5 Observations and Recommendations

In the preceding sections, the Panel identified key research themes and goals, as well as specific recommendations for the medical informatics field in the years ahead. It is evident that there is much overlap among the issues and strategies that were the focus of those sections. In this chapter, the Panel summarizes overall recommendations that follow from a synthesis of the topics discussed. The Panel believes that development of the medical information capabilities it has outlined is attainable through a national research program. The Library is a natural and appropriate locus for leading and sponsoring such an effort. It will lead other libraries through an evolutionary change from passive storehouses to partners in knowledge management.

Although the scenario in section three presented a 20-year goal, the Panel believes that the detailed mechanisms for reaching that level of performance are best expressed in terms of more short-term goals, albeit ones that are clearly aimed at the more distant goals outlined earlier. Thus, this chapter identifies specific goals not for 2006 but for 1996. It then describes the means to attain them and suggests desirable roles for the Library to adopt or maintain in support of those means.

The Panel has attempted to relate these recommendations to the financial resources needed to carry them out. These are expressed in terms of projects or activities, with dollar amounts roughly based on fiscal year 1985 costs. Intramural NLM costs were difficult to estimate; here, the Panel tried to express its resource estimates in terms of the time and effort required.

Institutional Responsibility

Goal

As a goal for 1996, NLM will become recognized as the principal national institutional locus for research and related matters that concern the organization and accessibility of biomedical information. This enhanced understanding of NLM and its role will reach all health-care professionals, biomedical researchers, NIH colleagues, and the public. NLM will accordingly be recognized and supported as the primary NIH institute for medical informatics research. It will be the NIH lead institute for agency-wide tasks and projects regarding medical knowledge issues (including both organization and distribution of such knowledge). It will serve DHHS and the health community as the coordinating headquarters of a unified system for medical terminology and hence for enhanced communication regarding medical data and knowledge.

Means to the Goal

As a means to achieving this goal, NLM must increase its intramural and extramural research. NLM's current authorization makes possible a full range of research and development activities, including support of demonstration projects. There is, therefore, no impediment to taking advantage of many opportunities, some of which may be in very applied domains. To remain in the vanguard, however, is a major institutional responsibility that requires a strong posture at the theoretical or *cutting edge* of the field. The field as a whole must be wary of confining itself to empirical problem solving, important though this is. To gain and enhance credibility as a scientific discipline, medical informatics workers must always keep the larger, more general considerations in mind. There must be a universal expectation of scientific rigor in projects, presentations, and papers of every kind, and medical informatics research should be reported not only in medical journals, but in general science publications as well.

Encouragement of young investigators and sponsorship of their professional development is an institutional responsibility that concerns NLM, professional medical informatics societies, and academic medicine. Sponsorship of training is a specific Panel recommendation outlined below, but related to it are many activities that give the field visibility and make it attractive to young talent.

To achieve major advances, there is a need for a consistently maintained infrastructure, both intellectual and technological. The intellectual infrastructure is the identification and communication of biomedical knowledge-management techniques and representation methods. The development, validation, and management of a unified system for encoding biomedical text and knowledge is a key example of such an activity. A unified language system will require an institutional base of credibility and coordination.

The technological infrastructure for such advances will be a computer-based communications network, as recommended in another section of this report. This will make possible a new dimension of collaborative research and collegial activity, not only for medical informatics but eventually for general information exchange in biomedicine. Once initiated, such a network must be maintained and administered.

NLM's role is to accept the responsibility for leadership. Formal identification as the 'National Institute for Medical Informatics' would be a highly desirable subtitle for the Library (at least one Panelist dissents from this recommendation, feeling it is tactically unwise at the present time). Leadership, however, must be shared; the well-accepted apparatus of peer group consultation must be expanded so that opportunities for consensus seeking are realized, and groups with related interests and objectives are both consulted and brought together. To arrive at consensus, achieve higher visibility, and identify new research opportunities and results, NLM should sponsor and support an ongoing series of conferences, colloquia, workshops, and symposia. Although NLM should vigorously exploit its own impressive laboratory and meeting facilities, sponsorship of similar meetings in other parts of the country should not be neglected. Joint sponsorship and cost sharing with other NIH agencies or foundations should be sought, in part to increase the circle of interest in the field and understanding of its scientific goals and content. The results of these meetings should be recorded and published in appropriate ways; innovative reports using video disk or videotape should be explored.

Carrying out these institutional responsibilities will inevitably involve a greater level of operational effort than at present. Up to three FTE positions may be needed. If NLM itself holds three meetings a year, internal operational funds will be required. Additional new funds will be needed to support publication of results, sponsorship of meetings (or activities) elsewhere, and support of joint activities with relevant professional societies or colleges.

Specific Recommendations

In order to assume responsibility while establishing itself as the national center for medical informatics activities, the Library should do the following:

(1) Seek a formally recognized *subtitle* as 'The National Institute for Medical Informatics,' to be reflected on NLM publications, letterhead, and other communications.

(2) Expand and maintain its Learning Demonstration Center for providing the public (and those appropriate parts of the Federal Government) with information regarding basic and applied activities in medical informatics research.

(3) Seek to interface broadly in an advisory capacity with other NIH institutes.

(4) Expand its activities in support of young investigators through its training programs, New Investigator Awards, and Research Career Development Awards.

(5) Introduce a formal in-house sabbatical program that would bring medical informatics investigators and trainees to NLM for periods of 3, 6, or 12 months.

(6) Hold two to three workshops and small conferences annually in targeted areas of high relevance to medical informatics.

(7) Sponsor additional workshops and small conferences through an external grants program that encourages joint support from other Federal agencies or private foundations.

(8) Establish and maintain a communications network for NLM grantees and others doing research in medical informatics, thereby gaining experience and insights that can in time lead to a national biomedical communications network.

Unified Medical Language System

Goal

The goal for the next decade is to attain professional acceptance by the national biomedical research and health-care communities of a standardized nosology and terminology for medical language, with a defined structure and syntax. By 1996, the elaboration of this Unified Medical Language System will be at the stage where numerous efforts to implement it in clinical settings are well underway, together with exploration of effective representation options.

Means to the Goal

Achieving this goal will require both coordination and ongoing basic research. The coordination task will emphasize creation of a highly committed medical community, organized into commissions and panels representing all the health professions. Cooperation and coordination will call for sensitive organizational leadership on the part of Federal programs, such as NLM, NIH, and DHHS, as well as national health organizations and major foundations. These efforts must be complemented by fundamental research. Research issues include better understanding of the inherent structure and organization of medical knowledge and information, building on the substantial insights already gained through NLM's work in developing the well-accepted and standardized MeSH system. Other issues include the development of effective computer-based representations of knowledge and enhanced understanding of the linguistic features in medical terminology. The research must be accompanied by the iterative development of terminologies and nomenclatures, with related development of electronic dictionaries and thesauri linked through computer-based indexing technologies.

NLM's role will include major responsibility for coordination and communication among the professional societies and institutions concerned. In cooperation with NIH, the Public Health Service, and relevant professional bodies, NLM should establish an oversight group to plan the development and eventual implementation of a unified and standardized medical language system. NLM should also support research in language attributes, terminology, representation, and refinement of computational tools for testing and implementing the results of these labors. Intramurally, the research effort should be a major program of high priority, directed in large part to interface with bibliographic systems and other activities in the library community. Extramurally, there should be defined linkages with centers of excellence in medical informatics.

Specific Recommendations

NLM should undertake the following activities in support of the development of a Unified Medical Language System:

(1) Serve as integrator for cross-professional activities to arrive at specifications for an acceptable unified language system.

(2) Establish high-priority intramural research activities in medical language, knowledge representation, and computer-based encoding schemes.

(3) Complement intramural research activities with targeted contracts to external research groups.

(4) Institute an extramural grants program in medical language system development to support approximately eight investigator-initiated basic research projects.

Communications Network

Goal

By the end of the next decade, there will be a national computer network for use by the entire biomedical community, both clinical and research professionals. The network will have advanced electronic mail features, as well as capabilities for large file transfer, remote computer log-in, and transmitted graphics protocols. It will either be part of a larger national network of scientists or will have gateways to other federally sponsored networks.

Means to the Goal

To achieve this goal, Federal sponsorship, direction, and funding are necessary. There are several possible approaches, keeping in mind the demanding specifications of the 10-year goal. Reliance on a simple electronic exchange with electronic bulletin boards will not suffice. Design and implementation of a totally new communication network would assure the most advanced system possible, but it would require a large amount of initial capital. Participation in an already established network would undoubtedly be more cost-effective, but such participation or cosponsorship would depend on successful negotiation among executive departments. Of the established networks, the ARPANET system most closely meets the needs of the health community, having all the capabilities described above. Although even that technology is aging and likely to change substantially in the next decade, it still serves as a useful model for the *kind* of performance that would be desirable. ARPANET is well accepted by its national scientific community, and its users already include artificial-intelligence-in-medicine investigators. The Department of Energy is one of its cosponsors.

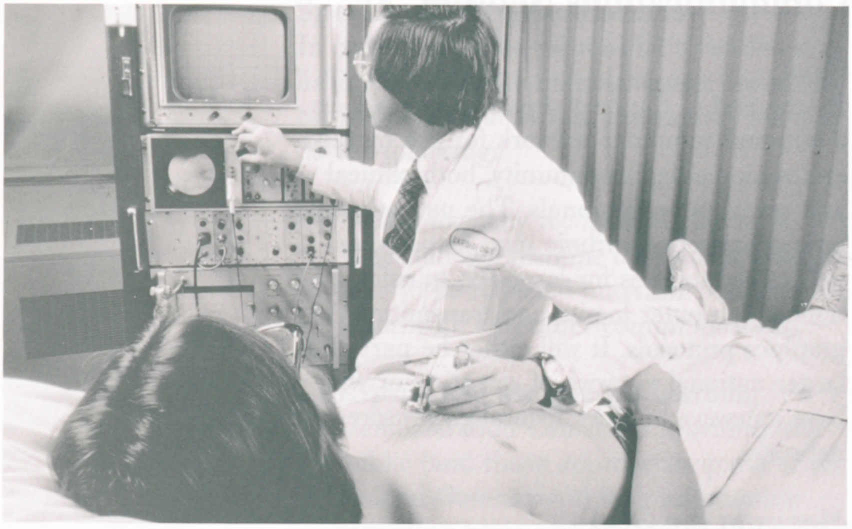

The Panel believes that a medical communications network should grow incrementally, with steady expansion based on improvement through experience. Biomedical research computers, as well as hospital- and clinic-based machines, should gradually be added to permit access via local nodes that meet local needs. Such growth will depend partly on demonstrating and publicizing the advantages of such facilities to the health professions and partly on the availability of technical support for local problem solving.

In describing a specific role for NLM, the Panel realizes that what is proposed is also an NIH issue and should be recognized as such from the beginning. NLM, however, should serve as the NIH lead institute for exploring prospects for cosponsorship or participation with DARPA. The network should begin on a small scale, starting with NLM's medical informatics awardees, advancing to the medical informatics community as a whole, to the academic medical library community, then to medical research investigators, and, eventually, to all health providers. NLM should also assume the responsibility for organizing user advisory groups, determining specific medical information protocols, and exchanging information models among users.

Specific Recommendations

To achieve the 10-year goal of a national biomedical communications network, NLM should do the following:

(1) Immediately begin to explore the communication options available via existing Federal and commercial networks.

(2) Establish a coordinated network for medical informatics researchers, using the mechanism selected in the evaluation phase, providing expertise to assist with joining the network as well as centralized coordination of the communications tasks involved.

(3) Move cautiously to increase the size of the network and the kinds of users served, establishing educational and public relations programs as appropriate.

(4) Work with the medical informatics community to establish standards and develop specialized software tools for knowledge and data exchange over such a network.

(5) Encourage costs for adding new machines to the network to be borne locally, but permit requests for grants to bring key machines online as appropriate.

Centers of Excellence

Goal

By the end of the next decade, there will be approximately 15 national centers of excellence in medical informatics, located at major academic medical centers. The emphasis in these centers will be on academic activities; the research carried out under their auspices will provide the organizing focus for defining the entire field. This research, rigorous in its scientific scholarship, will be coupled with numerous research training activities.

Means to the Goal

The establishment of productive, well-supported centers of excellence conducting basic research and training in the fundamental problems of medical informatics is crucial to the growth of the field as a respected, contributing scientific discipline. This goal can be achieved only through a commitment of national resources over a sustained period of time. The growth of the field and its potential contributions have been impeded by past uncertainties of support. The field has also suffered from unrealistic expectations about the length of time needed to produce and demonstrate research results. Center grants would provide stable sources of basic support for top research and training institutions, while permitting encouragement of focused centers in subareas of the medical informatics field. Five-year center grants would be supplemented by investigator-initiated grants supporting specific research projects.

NLM's role should be to accord high priority to support of such centers through its extramural research program. The center grant, a funding mechanism used by other NIH institutes, would be an appropriate category of award to achieve these goals. NLM should also provide, concurrently, innovative opportunities for exchange efforts among the centers it supports and its own intramural laboratories. NLM's Lister Hill National Center for Biomedical Communications should, in fact, function as such a center and participate fully as a collegial member of this emerging academic community.

Specific Recommendations

In order to encourage the development of centers of excellence in medical informatics, the NLM should do the following:

(1) Initiate six such centers as soon as possible through a competitive extramural grants program that would assure a minimum of five years' support to grantees.

(2) Encourage ongoing applications for center grants until there are 15 active centers at the end of 10 years.

Training in Medical Informatics

Goal

By 1996 one fifth of American medical schools will offer formal career training in medical informatics, and all American health science schools will offer training in medical informatics as part of their curricula. Careers in medical informatics will offer many opportunities. There will be a growing need for research investigators in academic settings, as well as in an evolving medical informatics industry. Qualified persons will be needed to teach other workers in the field, as well as health professionals, to administer academic departments and carry out informatics services in health-care organizations of every kind.

Means to the Goal

As a means to this ambitious goal, a first step will be to define better the range and nature of desirable educational programs. Questions of curricular content, appropriate degree (if any) to be earned by trainees with different career goals, professional level (pre- or postdoctoral), and job training standards all need to be addressed, perhaps with alternative approaches being tried in experimental or pilot programs. The importance of this field as a medical discipline must be brought home to leaders in academic medicine. Centers of excellence must be expected to emphasize training and cultivate ties with technical and graduate schools as appropriate. Professional societies and foundations can make a large contribution, particularly valuable in the near future, by supporting special informatics fellowships for clinicians on sabbatical or other leave.

NLM's role is to continue its successful grant program for research career training in medical informatics, increasing the number of trainees per institution to a minimum of six, increasing the number of institutions to 10 as soon as possible, and adding one more program each year over the next 10 years. In directing the growth of its grant awards, NLM should consider, with advice of appropriate consultants, what modifications or new directions might be usefully adopted. The training offered, however, should be academic in character; support for vocational training should be sought from other sources. Training ought also to be offered at NLM on a short-term basis or as an in-house sabbatical for an academic year.

Specific Recommendations

To enrich the quality and impact of its support for medical informatics training, NLM should do the following:

(1) Increase immediately the number of NLM-funded training programs, each supporting a minimum of six trainees, from 5 to 10 institutions.

(2) Gradually increase the number of training grants, by an increment of approximately one per year, until 20 institutions are receiving NLM support at the end of the next decade.

(3) Establish stipends to support five in-house sabbaticals annually at NLM for mid-career professionals.

Cognition and Decision Support

Goal

By 1996, commercially marketed decision-support systems for limited, narrow domains of medicine will be generally accepted by professional users and widely available. Methods for broadening their coverage and impact will be generally recognized and close to the implementation stage. Developing systems will use information displays that are sensitive to cognitive principles that have been delineated through ongoing research regarding human-machine interaction.

Curricula and textbooks for teaching medical decision making will be available in health professions schools; courses will exist at all institutions. In medical schools, decision-making principles will be integrated into regular clerkships. A working knowledge of principles of decision making will be presumed to be part of what it means to be a health professional. There will be enough qualified, trained medical decision-science experts to teach these courses at most professional schools on an advanced level. Ongoing research in the cognitive disciplines, such as decision making, will be underway in numerous academic medical centers, and certainly in all those held in esteem for the quantity and quality of their biomedical research.

Means to the Goal

As a means to this goal, research, particularly on fundamental issues, needs far more emphasis and support than it receives at present. In addition to research, orientation to decision sciences and related cognitive disciplines must be offered to clinical faculty as part of their continuing professional development.

The research agenda in this area is substantial. There is a need for more work on cognitive processes, particularly a better understanding of problem formulation. There is a similar need for work on values, preferences, and utilities. Mathematical modeling techniques also require urgent attention. The function of nonverbal information in decision making is not well understood and warrants study. Models of system users and of the nature of human-machine interaction are needed, particularly for optimizing effective displays of medical information and developing user tools for individual workstations. There is also a continuing need for descriptive studies of normative behaviors in all the health professions. The organization and use of knowledge from diverse sources must also be better understood. Such research provides important background for investigation in other areas of medical informatics, such as system evaluation and the development of a Unified Medical Language System. The integration of such concepts in a coherent research agenda is most likely to occur in the centers of excellence, as recommended earlier.

Faculty development is a significant concern if decision-making training is to become a standard element in the formation of health professionals. One mechanism is to make available a wide range of orientation and training activities. These could range from postdoctoral training programs for clinical faculty to sabbaticals at informatics centers of excellence, summer residencies or seminars, short-term workshops, and sessions at national or regional meetings of professional societies. There is a particularly good opportunity here for significant contributions by professional societies and foundations.

NLM's role in this area should be to support and encourage both research and faculty development. In doing so, NLM should seek and stress cooperation and coordination with other appropriate sources, such as foundations and societies. Some of these activities overlap with interests of other programs and agencies, so possibilities of joint support ought to be sought. Those programs whose aim is the improvement or provision of biomedical research resources also ought to become involved. In the research it supports, NLM should emphasize basic issues and fundamental work, but with testing in real environments. Awardees should be expected to make contributions to science, not just adapt an existing technique to a new problem.

Specific Recommendations

To achieve the goals for progress in cognitive science and decision support outlined above, NLM should do the following:

(1) Immediately increase the extramural support for research in this area to a minimum of eight researcher-initiated grants per year, with an additional 50 percent increment annually for five years.

(2) Encourage research activities in this area through targeted programs mentioned in earlier sections: centers of excellence, NLM-sponsored workshops, New Investigator Awards, and Research Career Development Awards.

(3) Actively encourage co-funding or research with other appropriate Federal agencies, particularly at NIH.

(4) Promote faculty development programs in medical decision making by sponsoring workshops, symposia, and NLM sabbaticals, particularly seeking joint funding or sponsorship with foundations or professional societies.

Knowledge Bases and Data Bases

Goal

By 1996, medical libraries will serve as gateways and promoters for many new forms of biomedical information exchange. The medical literature will largely be available electronically in full text, with print-on-paper as an option. The literature will also be available universally and rapidly, with easy accessibility at nearly every point of health care. "Smart" retrieval systems will exist, at least in prototype form.

A wide variety of clinical and research data bases will be used increasingly as knowledge sources. Validated methods, in prototype form, will exist for extracting knowledge from such data bases semi-automatically. Electronic networks, which permit the functional distribution of work as a joint exercise among geographically distant persons and institutions, will be used routinely for knowledge-base organization, validation, and maintenance. These knowledge bases will form an essential software component for increasingly important decision-support systems. Integrated knowledge bases will include nonverbal as well as verbal information. They will use sophisticated technology for handling pictures and other signals and will often be able to store, organize, and retrieve not only the images or signals themselves, but also their informational content. Laboratory scientists will depend on large national data bases for management of large bodies of research data and analysis capabilities that could not realistically be provided locally (e.g., gene sequence data bases for use by molecular biologists).

Means to the Goal

As a means toward these goals, research is needed on a variety of representation issues for all biomedical knowledge. Significant issues that are beginning to be addressed include the representation of time-dependent phenomena, the difficulty of describing process in data bases, and the role of taxonomies and classifications as they affect medical language. More must be known about the context of dependent circumstances in representation of biomedical knowledge of every kind. Automatic indexing systems must be developed and made cost-effective; efficient knowledge base editing and updating methods must be perfected, and useful tools developed.

NLM's role, as the central institution for medical informatics research, is to sponsor this kind of research in its extramural programs. When appropriate, NLM should also develop or contribute to specialized research data bases, such as the GenBank system already in use by the molecular biology community. In NLM's intramural activities, there should be a concentration on the language, full-text representation, and indexing research activities now well underway. NLM and its Regional Medical Libraries should continue to serve as major nodes in a larger medical communications network (see discussion of communications network above). Sabbaticals at NLM for health professions faculty should be encouraged because of the unique laboratory aspect of its indexing and retrieval experiments. In these ways, NLM can perform an invaluable service by promoting the unification and sharing of medical knowledge bases.

Specific Recommendations

In order to promote the development of innovative knowledge bases and data bases of use to health professionals and biomedical researchers, NLM should do the following:

(1) Immediately double the extramural support for research in this area to approximately eight researcher-initiated grants, with an additional 50 percent increment annually for five years.

(2) Encourage research activities in this area through targeted programs mentioned in earlier sections: centers of excellence, NLM-sponsored workshops, New Investigator Awards, and Research Career Development Awards.

(3) Actively encourage co-funding or research with other appropriate Federal agencies, particularly at NIH. This involves analyzing current NIH and NLM programs to identify those that are highly appropriate for joint funding in terms of major positive impact in both the field of primary study and the medical informatics field.

(4) Continue an active intramural research program in the areas cited, but increase the investment by at least 50 percent in the first year.

(5) Expand the activities of the Regional Medical Libraries in coordinating the development and linkages of national shared knowledge and data bases.

Supporting Technologies

Goal

By 1996, there will be generally available specialty-oriented computer workstations for those in many professional specialties and research disciplines. As a rule, these workstations will be integrated via networks with hospital information systems, and many of these systems and workstations will provide decision-support capabilities. Voice query systems will be readily available and easy to use. Complexities of operating systems, network connections, and related computer science technologies, at present so bedeviling to the naive amateur, should become nearly transparent to the health professional.

Means to the Goal

The means by which these supporting technologies will be significantly advanced or perfected will be a variety of research and development activities now being pursued on a large scale in academic institutions and industry. Industrial development is currently oriented towards general scientific workstations; thus, cultivating the establishment of a medical informatics industry will depend on the perceived usefulness and desirability of specialized software or hardware tools on the part of both health professionals who use such systems and their employers.

NLM's role should be to support the medically significant portion of this developmental work as part of its research programs. The task is much too large for NLM to attempt to assume it all. There should be a continuing willingness to support work where the relevance of these technologies to biomedical domains can be explored or exploited. For example, a specialized workstation for medical librarians would be a particularly appropriate intramural development project. Features might include decision-support software to assist with cataloging, or specialized devices for accessing diverse, geographically distributed data bases and teaching their use to health personnel.

It will be important to cultivate and maintain ties with academic computer science departments, as well as with industrial concerns. At least some of the centers of excellence that NLM will support should become known for these computer science/industrial connections. These connections should take place in a research context; examples of potential benefit are in the areas of enhanced information storage devices and improved interfaces for graphical images and knowledge-base examination. Related exploratory work is especially appropriate for NLM's intramural laboratory activities. The amount of activity should probably remain at current levels, with the exception of a new effort to develop a customized workstation for use by medical librarians.

Specific Recommendations

To contribute to the development of supporting technologies customized for use by biomedical professionals, NLM should do the following:

(1) Continue the intramural development of demonstration projects in specialized informatics hardware, adding as a specific task the development of software, and hardware if needed, for a specialized workstation for medical librarians.

(2) Announce a willingness to support innovative research in the use of new informatics technologies to support specific biomedical needs, and fund three to five awards annually in this area (in addition to such work being undertaken in the centers of excellence).

Evaluation

Goal

At present, informatics research is both helped and hampered by problems of evaluation design. The design of formal evaluation has its own disciplinary apparatus and methodology, but incorporating these techniques for the validation of informatics technologies has often been difficult, both conceptually and in practice. Particularly difficult has been the determination of acceptable levels of performance in areas where there is often disagreement among professionals, even among experts with similar training and experience. Within 10 years, there will be a more adequate theory and technique, accepted by the medical informatics community and health science professionals. It is especially important that appropriate evaluation guidelines be developed for the validation of medical knowledge bases and decision-support systems.

Means to the Goal

As a means to this end, multidisciplinary research should be pursued that will involve full collaboration and participation of the medical informatics community with social and behavioral scientists, as well as with domain experts in clinical specialities and health professions. The research must be augmented by the focused attention of professional societies and academic medical institutions. There should be consensus as to when evaluation is needed and appropriate, where it should fit in advanced informatics research, and how it best responds to practitioners' straightforward but vital concerns with safety and usefulness.

NLM's role should be to keep the issues of evaluation design visible and support some research into theory and methodology. NLM's consultants, reviewers, and program staffs should constantly be reminded that a major issue is to determine when evaluation is relevant and helpful, and to avoid making sophisticated evaluation design a requirement when such efforts would be premature. Sponsorship of occasional consensus conferences on the status of evaluation methodology relevant to medical informatics would benefit all concerned. An appropriate topic for a specialized workshop, for example, would be the exploration of avenues to provide development grants or other mechanisms that could help carry the products of medical informatics research closer to commercial viability.

Specific Recommendations

To encourage the emergence of improved methods for evaluating the products of medical informatics research, as well as the development of well-validated systems themselves, NLM should do the following:

(1) Support three or four investigator-initiated projects in the area of evaluation research.

(2) Encourage in its extramural grants program the inclusion of appropriate evaluation studies (*when such validation experiments are appropriate and not a distraction from the basic research activities that often precede the development of working prototypes suitable for study*).

(3) Sponsor workshops and consensus conferences in the area of medical informatics technology evaluation and transfer, ideally achieving joint sponsorship with other agencies that have an interest in technology assessment and dissemination.

6 Priorities and Summary

Priorities on the Recommendations

After the Panel had developed the specific recommendations outlined in the previous chapter, it sought to develop priorities to guide the implementation process. This type of prioritization was difficult for two reasons. First, there is substantial interdependency among the recommended actions, and these relationships make it difficult to identify any one direction to pursue at the expense of the others. Second, and related to the first, there was general recognition of the need for a broad mandate, one that acknowledges the need for the full range of tools—centers of excellence, grants, training, workshops, communications, etc. The future of the research arm of the Library, and indeed of medical informatics as a discipline in the United States, is intimately tied to the full complement of activities and the broad base of support that has been proposed.

With these caveats, we present here the relative priorities on the items outlined in section five. The recommendations fall into five priority groupings, but it should be emphasized that the item in the fifth group is by no means of low priority—it simply came last when the Panelists were forced to choose.

The priority groups are as follows:

(1) *Institutional Responsibility:* It was uniformly agreed that the other recommendations hinge in large part on NLM successfully establishing itself as the national focus for medical informatics activities. The dollar amounts required are relatively low, but the potential for influence is great. The specific recommendations outlined under this item in section five are given the greatest priority.

(2) There are three major categories of activity that fall in the second priority group:

- *Centers of excellence*
- *Medical informatics training*
- *Unified Medical Language System*

(3) The third priority group includes another three categories of recommendation:

- *Communications network:* This was the only category for which the prioritization process showed disagreement among the Panelists. Several individuals, generally those who have had extensive personal experience with electronic communication networks, felt that NLM's role in promoting an integrated network for the biomedical community was of especially high priority as an enabling activity for all the other recommendations.
- *Cognition and decision support*
- *Knowledge bases and data bases*
- *Evaluation*
- *Supporting technologies*

With this list providing a general sense of priorities for proceeding with the ambitious research and development agenda proposed, we summarize in the next section the Panel's recommendations for developing an infrastructure and then promoting innovative research.

Summary of the Recommendations

The Planning Panel on Medical Informatics Research has been guided by an effort to anticipate future health-care and research needs and the role that information and communication technologies can play in responding to them. It has developed a scenario for the future (section three), parts of which may be achieved in 20 years or so if research and logistical resources are properly applied. It has reviewed the state of the art in medical informatics research (section four) and has proposed strategies for guiding current research toward the future goals it has identified. Finally, it has made specific recommendations to the Library (section five), identifying how that agency can catalyze the further emergence of a medical informatics discipline and help disseminate its research results in the years ahead.

We summarize the Panel's recommendations here in two categories: (1) enabling mechanisms for improving research productivity and bringing its results to the biomedical research and health-care communities, and (2) research topics that must be supported and pursued if our 20-year view of the future is to be achieved.

Enabling and Logistical Activities

Medical informatics research has already made enormous strides, but the problems of technology transfer continue to hinder its full impact in the health-care setting. NLM is encouraged to work with the research community to develop the infrastructure needed for more effective sharing of research results, their improved evaluation, and their facilitated dissemination. Specific issues include:

- Identification of NLM as the central Government organization for the medical informatics field, one that works closely with the rest of NIH and complements the information dissemination and knowledge-management activities of the other institutes;

- Sponsorship of workshops and conferences, especially in areas that are hindering the dissemination of 1980's technology (e.g., evaluation methodologies, legal barriers, lack of data, and terminology standards);

- Establishment of a national communications network, initially for NLM grantees, then for biomedical researchers, and ultimately for the entire biomedical community;

- Creation of centers of excellence in medical informatics research and training;

- Support of training programs in medical informatics, complemented by training opportunities at NLM itself and strong programs in Research Career Development and New Investigator Awards; and

- Continuation of strong intramural research programs in medical informatics and support for investigator-initiated grants in the field.

Research Priorities

Research priorities in medical informatics are guided by our need to acquire, structure, provide access to, and teach the knowledge of medicine and health.

- We must *accumulate* and *synthesize* medical expertise so that it is easily accessible to the next generation, as well as to individuals in the present generation. This requires an understanding of the semantic structure in data and knowledge and the development of effective methods for encoding medical information. Our enthusiasm to develop a unified supporting system for medical language must be tempered by the realization that tomorrow's knowledge may not mesh well with today's framework; novel techniques to assure flexibility are therefore mandatory.

- We must *deliver* knowledge and relevant data, providing them more reliably, more cheaply, and with greater selectivity than is possible with present methods (books, journals, consultants, and today's computer-based resources).

- We must *teach* the knowledge of medicine and health in new ways that are both sensitive to the limitations of human memory and traditional classroom approaches and guided by emerging technologies for instruction and self-study.

These goals require focused research in several key areas of medical informatics. High priority is placed on research in cognition and decision support, on the development and use of knowledge and data bases, on the creation of a Unified Medical Language System, and on the development of new methodologies for evaluating the results of medical informatics research.

Conclusion

It is a humbling experience to try to anticipate the future as we have done in this report. The world is radically different from what any of us would have predicted two decades ago, and there is little reason to believe we will be any more correct in trying to anticipate what lies 20 years in the future. Yet it is in the attempt to do so that progress is made, for views of what *ought* to be or what *might* be are sources of inspiration and can subsequently affect decisions regarding resource commitment. The program proposed here has resulted from an effort to anticipate what *could* and *should* be done; we encourage the commitment of resources that will allow progress in medical informatics to take its course. The health care of the Nation has much to benefit when the stated goals are achieved.

References

1. Association of American Medical Colleges. *Report of the Steering Committee on the Evaluation of Medical Information Science in Medical Education.* Washington, 1986.
2. Ziman JM. Information, communication, and knowledge. *Nature* 1969;224:318-24.
3. de Solla Price D. *Little Science, Big Science.* New York; Columbia University Press, 1963.
4. Williams ME. Transparent information systems through gateways, frontends, intermediaries, and interfaces. J AM Soc Infor Sci 1986; 37:204-14.
5. Pauker SG, Gorry EA, Kassirer JP, Schwartz WB. Towards the simulation of clinical cognition: taking a present illness by computer. *Amer J Med* 1976;60:981-96.
6. Stross JK, Harlan WR. The dissemination of new medical information. *JAMA* 1979;241:2622-24.
7. Schaffner W, Ray WA, Federspiel CF. Surveillance of antibiotics in office practice. *Ann Intern Med* 1978;89:796-9.
8. Holden C. An omnifarious data bank for biology? *Science* 1985;228:1412-13.
9. Committee on Models for Biomedical Research. *Models for biomedical research—a new perspective.* Washington DC: National Academy Press, 1985.
10. Bush V. As we may think. *Atlantic Monthly* 1945; 176(1):101-8.
11. Licklider JCR. *Libraries of the Future.* Cambridge, MA: MIT Press, 1965.
12. Warren K, ed. *Coping with the Biomedical Literature Explosion: A Qualitative Approach.* New York: Rockefeller Foundation, 1978.
13. White K. Information for health care: an epidemiological perspective. *Inquiry* 1980; 17:296-312.
14. Lindsay RK, Buchanan BG, Feigenbaum EA, Lederberg J. *Applications of artificial intelligence for organic chemistry—the DENDRAL project.* New York: McGraw-Hill, 1980.
15. Shortliffe EH. *Computer-based medical consultations: MYCIN.* New York: American Elsevier Computer Science Library, 1976.
16. Miller RA, Pople HE, Myers JD. INTERNIST-1: an experimental computer-based diagnostic consultant for general internal medicine. *N Engl J Med* 1982; 307:468-76.
17. Feigenbaum EA, Barr A, Cohen PR. *The handbook of artificial intelligence*—volumes 1, 2, and 3. Los Altos, CA: Kaufmann, 1982.
18. Brachman RJ Levesque HJ, eds. *Readings in knowledge representation.* Los Altos, CA: Morgan Kaufmann, 1985.
19. Lenat D, Prakash M, Shepherd M. CYC: using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI Magazine* 1986; 6 (4): 65-85.
20. Committee on Models for Biomedical Research. *Models for biomedical research—a new perspective.* Washington DC: National Academy Press, 1985.
21. Buchanan BG, Shortliffe EH. *Rule-based expert systems.* Reading, MA: Addison-Wesley, 1985.
22. Pople HE. Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnostics. In: P. Szolovits, ed. *Artificial intelligence in medicine.* Boulder: Westview Press, 1982.
23. Miller PL. Attending: critiquing a physician's management plan. In: *IEEE transactions on pattern analysis and machine intelligence*, vol. 5. Washington, DC: IEEE Computer Society Press, 1983: 449-61.
24. Patil RS. *Causal representation of patient illness for electrolyte and acid-base diagnosis.* Ph.D. dissertation, Massachusetts Institute of Technology, MIT/LCS-TR 267, 1981.
25. Pauker SG, Kassirer JP. Therapeutic decision making: a cost-benefit analysis. *N Engl J Med* 1975; 293:229-34.
26. Shortliffe EH, Scott AC, Bischoff M, Campbell AB, van Melle W, Jacobs C. ONCOCIN: an expert system for oncology protocol management. In: *Proceedings of the seventh international joint conference on artificial intelligence.* Vancouver, BC: 1981:876-81.
27. Rosati RA, McNeer JF, Starmer CF, Mittler BS, Morris JJ, Wallace AG. A new information system for medical practice. *Arch Intern Med* 1975; 135:1017-24.
28. Blum RL. *Discovery and representation of causal relationships from a large time-oriented clinical data base.* Ph.D. dissertation, Computer Science Department, Stanford University, 1982.

29. Masarie FE, Miller RA, First MB, Myers JD. An electronic textbook of medicine. In: *Proceedings of the ninth annual symposium on computer applications in medical care.* Washington DC: IEEE Computer Society Press, 1985: 335.

30. Davis R. *Applications of meta-level knowledge to the construction, maintenance, and use of large knowledge bases.* Ph.D. dissertation, Computer Science Department, Stanford University, 1976.

31. Politakis P, Weiss SM. Using empirical analysis to refine expert system knowledge bases. *Artificial Intelligence* 1984;22:23-48.

32. Lenat D, Prakash M, Shepherd M. CYC: using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI Magazine* 1986;6(4):65-85.

33. Michalski RS, Carbonell JG, Mitchell TM. *Machine learning: an artificial intelligence approach.* Palo Alto: Tioga Publishing Company, 1983.

34. Miller RA, McNeil MA, Challinor SM, Masorie FE, Myers JD. The INTERNIST-1/QUICK MEDICAL REFERENCE project-status report. *West J Med* 1986; 145:816-822.

35. Elstein AS, Shulman LS, Sprafka SA. *Medical problem solving: an analysis of clinical reasoning.* Cambridge MA: Harvard, 1978.

36. Kassirer JP, Kuipers BJ, Gorry GA. Toward a theory of clinical expertise. *Am J Med* 1982; 73:251-59.

37. Johnson PE. What kind of expert should a system be? *J Med Philosophy* 1983;8:77-97.

38. Elstein AS, Bordage G. The psychology of clinical reasoning. In: Stone G, Cohen F, Adler N, eds. *Health psychology.* San Francisco: Jossey-Bass, 1979.

39. Edwards W. Conservatism in human information processing. In: Kleinmuntz B, ed. *Formal representation of human judgment.* New York: Wiley, 1968.

40. Casscells W, Schoenberger A, Graboys TB. Interpretation by physicians of clinical laboratory results. *N Engl J Med* 1978;299:999-1001.

41. Kahneman D, Slovic P, Tversky A, eds. *Judgment under uncertainty: heuristics and biases.* New York: Cambridge Univ Press, 1982.

42. McNeil BJ, Pauker SG, Sox HC, Tversky A. On the elicitation of preferences of alternative therapies. *N Engl J Med* 1982;306:1259-62.

43. Cebul R, Beck L, Carroll G. *Teaching clinical decision making.* New York, NY: Praeger, 1985.

44. Weinstein MC, Fineberg HV. *Clinical decision analysis.* Philadelphia: Saunders, 1980.

45. Politser PE. Decision analysis and clinical judgment: a re-evaluation. *Medical Decision Making* 1981;1:361-89.

46. Fallside F, Woods W, eds. *Computer speech processing.* Englewood Cliffs NJ: Prentice-Hall, 1985.

47. Harris L, ed. The advantages of natural language programming. In: Sime, Coombs, eds., *Designing for human-computer communication.* London: Academic Press, 1983.

48. Thomson C. Using a menu-based natural language interface to ask spatial data base queries. In: *Proceedings of the 1984 Pecora conference on spatial information technologies for remote sensing.* Sioux Falls SD: 1984.

49. Davis R. Interactive transfer of expertise. *Artificial Intelligence* 1979;12(2):121-57.

50. Swartout W. XPLAIN: A system for creating and explaining expert consulting systems. *Artificial Intelligence* 1985;21(3):285-325.

51. McKeown K. *Text generation.* Cambridge, England: Cambridge Univ. Press, 1985.

52. Sleeman D, Appelt D, Konolige K, Rich E, Sridharan NS, Swartout W. User modelling panel. In: *Proceedings of the ninth international joint conference on artificial intelligence.* Los Altos, CA: Morgan Kaufmann 1985;1298-1302.

53. Richer M, Clancey W. GUIDON-WATCH: a graphic interface for viewing a knowledge-based system. *IEEE Computer Graphics and Applications,* 1985;51-64.

54. Goldwasser S, Reynolds A, Bapty T, Baraff D, Summers J, Talton D, Walsh E. A 3-D physician's workstation with real-time performance. *IEEE Computer Graphics and Applications,* 1985.

55. Chen L, Herman G, Hung H, Lefkovitz H, Trevedi S, Udupa J. Interactive manipulation of 3-D data via a 2-D display device. *Optical Engineering* 1985;24:893-900.

56. Cassell EJ. *Talking with patients.* Cambridge MA: MIT Press, 1985.

57. Jenkin MA. Clinical informatics: a strategy for the use of information in the clinical setting. *Med Inform 1984;*9:225-32.

58. Gabrieli ER. The medicine-compatible computer: a challenge for medical informatics. *Med Inform* 1984; 9:233-50.

59. Blum BI, ed. *Information systems for patient care.* New York: Springer-Verlag, 1984.

60. Ryan SA. Applications of a nursing knowledge based system for nursing practice: in service, continuing education, and standards of care. In: *Proceedings of the seventh annual symposium on computer applications in medical care.* New York: IEEE, 1983:491-4.

61. Evans S. A computer-based nursing diagnosis consultant. In: *Proceedings of the eighth annual symposium on computer applications in medical care.* New York: IEEE, 1984:658-61

62. Reggia JA, Tuhrim S, eds. *Computer-assisted medical decision making.* New York: Springer-Verlag, 1985.

63. Schulman JA, Crawford S, eds. Electronic publishing and health sciences libraries. *Bull Med Lib Assoc* 1986;74:19-40.

64. Association of academic health sciences library directors/Medical library association, Joint task force to develop guidelines for academic health sciences libraries. *Challenge to action: planning and evaluation guidelines for academic health sciences libraries.* Final draft, 1985.

65. Schoolman HM. The role of the medical library in support of clinical decision-making. Presentation at the symposium on medical education in the information age, Association of American Medical Colleges, 1985.

66. Abelson PH, Dorfman M, eds. Computers and electronics. *Science* 1982; 215:749-873.

67. Birnbaum JS. Toward the domestication of microelectronics. *Communications of the ACM* 1985; 28:1225-35.

68. Kleinrock L. Distributed systems. *Communications of the ACM* 1985; 28:1200-13.

69. Lundstrom SF, Larsen RL. Computer and information technology in the year 2000—a projection. *Computer* 1985; 18:69-79.

70. Bell CG. The mini and micro industries. *IEEE Computer* 1984;17:14-32.

71. Bell CG. Multis: a new class of multiprocessor computers. *Science* 1985:228;462-67.

72. Gerola H, Gomory RE. Computers in science and technology: early indications. *Science* 1984; 225:11-18.

73. Pohm AV. High-speed memory systems. *IEEE Computer* 1984;17:162-72.

74. Abelson PH, ed. Computers and research. *Science* 1985;228:401-70.

75. Lederberg J. Digital communications and the conduct of science: the new literacy. In: *Proceedings of the IEEE 1978;*66:1314-19.

76. Jennings DM, Landweber LH, Fuchs IH, Farber DJ, Adrion WR. Computer networking for scientists. *Science* 1986;231:943-50.

77. Hobbs LC. Printing and storage peripherals. *IEEE Computer* 1984:17:225-41.

Appendix A: NLM Planning Process

In January, 1985 the Board of Regents of the National Library of Medicine resolved to develop a long range plan to guide the Library in wisely using its human, physical, and financial resources to fulfill its mission. The Board recognized the need for a well-formulated plan because of rapidly evolving information technology, continued growth in the literature of biomedicine, and the need to make informed choices of intermediate objectives that would lead NLM toward its strategic, long range goals. Not only would a good plan generate goals and checkpoints for management, actually a map of program directions, but it would also inform the various constituencies among the Library's users about the future it sought and could help to enlist their support in achieving that future.

At the Board's direction, a broadly based process was begun involving the participation of librarians, physicians, nurses, and other health professionals; biomedical scientists; computer scientists; and others whose interests are intertwined with the Library's. A total of 77 experts in various fields accepted invitations to serve on one of the five planning panels. Each panel addressed the future in one of the five domains that encompass NLM's current programs and activities. The domains, which provided the panels, a framework for thinking about the future are:

1. Building and organizing the Library's collection
2. Locating and gaining access to medical and scientific literature
3. Obtaining factual information from data bases
4. Medical informatics
5. Assisting health professions education through information technology

The Library chose a planning model with three components. First, it incorporates a general, somewhat indistinct vision of the future 20 years from now in medicine, library and information science, and computer-communications technology. That environment cannot be forecast precisely, but we can speak of a "distant" goal. That goal is seen as a societal objective whose attainment involves many organizations and agencies. NLM has a major role to play in achieving the goal and must plan its part. Second, while the 20-year goals are indistinct, there are opportunities for and impediments to achieving them. The opportunities and impediments can be more clearly envisioned because they appear to lie roughly 10 years away. Third, the specific steps that should be taken to remove the impediments and take advantage of the opportunities should be programmed for 3 to 5 years.

The planning process also involved participation within the Library. The Director provided his version of the future in the form of a "Scenario: 2005," which was distributed to panel members and Library staff. NLM staff prepared background documents that reported NLM achievements in the five domains, and reviewed current planning. Senior NLM staff members also acted as resource persons to the planning panels.

At the end of the planning process, each panel formulated recommendations and priorities for future NLM programs and activities in the domain under its purview. The five panel reports were reviewed by the Board of Regents in June 1986. The Board then asked the NLM staff to analyze and reconcile their findings, eliminating any duplications and consolidating the recommendations. Together with the planning panel reports, this synthesized plan presents the official Long Range Plan of the Board of Regents of the National Library of Medicine.

Photographs were obtained from the several Bureaus, Institutes, and Divisions of the National Institutes of Health (including the Office of the Director, NIH, the Warren G. Magnuson Clinical Center, and the National Institute on Aging), the Uniformed Services University of the Health Sciences, the World Health Organization, and William A. Yasnoff, M.D., Ph. D..

December 1986