Computer Use in Diagnosis, Prognosis, and Therapy

Computers are used to simulate the reasoning of expert clinicians.

Harold M. Schoolman and Lionel M. Bernstein

Diagnosis, prognosis, and therapy are the three inextricably bound basic components of medical practice. We will review some of the major approaches of automated information management to these components. Computer systems for literature searches, hospital inventory control, housekeeping, billing, patient medication and procedure scheduling, patient monitoring, laboratory tests, diagnostic radiology, dosimetry in x-ray therapy, electrocardiogram interpretation, and so forth are well developed and widely used, but are not discussed here.

Summary. Computers are used to influence diagnostic and therapeutic decisions. The computer's information-handling capabilities allow it to serve as a reliable extension of the physician's memory and expander of the physician's information and synthesized knowledge resources. Computers have been used to facilitate decisions through organization of patient data, improved classification of patients, decision analysis in clinical settings, and simulation of expert clinical reasoning. Computer programs are more successful in narrow, constrained, single arenas of medicine with much underlying pathophysiologic understanding and where decisions are based largely on hard laboratory data. New models of synthetic reasoning that simulate expert clinical behavior show promise of supporting complicated decisions concerning problems of multiple diseases. All systems are confronted by problems of consensus and authority of the underlying information used.

Although diagnosis, prognosis, and therapy are better today than at any time in the past, there is now a greater dissatisfaction with them than ever before. Some suggest that the problem results from faults in basic medical education. Their argument is that there are unrecognized and untaught distinctions among the acquisition of information, the synthesis of information (knowledge), the use of information for problem-solving (competence), and the performance of physicians in actual patient encounters (*1*). Medical education has in the past, by concentrating on the first two steps, stopped short of its critical objective. Today "problem-solving" medical curricula are being developed. The problem-solving curriculum is usually described in terms of competence as measured by ability to solve test problems on paper rather than in actual performance, but there is no doubt that performance is the ultimate objective.

The goal is straightforward: to maximize the likelihood of "correct" decisions about diagnosis, prognosis, and therapy. The decisions depend on the physician's fund of information and ability to synthesize that information and appropriately apply it to problem-solving. The rapid growth of biomedical information has created an available body of knowledge (facts, concepts, and their interrelationships) far greater than any individual can assimilate. Even the most industrious and astute physician uses less than the total amount of potentially available relevant information and knowledge in making decisions. Computers, with their massive information-handling capabilities, are looked to as potential expanders of the physician's information and knowledge resources. It is in this restricted sense that the role of computerized medical information systems used in diagnosis, prognosis, and treatment will be reviewed by describing examples of some of the various approaches.

Computer-Based Consultation

Bleich (*2*) has constructed a computer program to help physicians manage patients with electrolyte and acid-base disorders. The program directs a dialogue in which the physician (or other user) enters clinical and laboratory information. The data are checked for proper syntax, internal consistency, compatibility with life, and so on. On the basis of abnormalities detected, the program asks additional questions to further characterize the electrolyte and acid-base disturbance. During or at the completion of the dialogue, an evaluation note is produced. As might a consultant's note, the program evaluation provides a list of diagnostic possibilities, an explanation of underlying pathophysiology, therapeutic recommendations, precautionary measures, suggestions for further studies, and references to the medical literature.

This program addresses an important area of medicine covering many patient encounters. The subject is one in which reliable laboratory tests provide much of the necessary information and about which considerable pathophysiologic understanding exists. Thus, the programming of definitive advice regarding decisions about individual patient care is greatly facilitated. To use the program successfully the physician must not only have identified the appropriateness of the subject for the patient but must also understand that the program does not take concomitant disorders into account. The system can readily accommodate additions or corrections, and its information data base has the authority of a consensus of a group of physicians and physiologists who have been in the forefront of research and clinical work in this subject. It provides information and is highly directive. The information provided significantly expands the working memory of facts and concepts.

The system is in wide use; although Bleich cautions (*3*) that "repeated usages cannot be taken as evidence of a positive contribution to patient care," it is probable that, in the area of electrolyte acid-

Dr. Schoolman is the deputy director for research and education and Dr. Bernstein is the assistant deputy director for research and education at the National Library of Medicine, Bethesda, Maryland 20014.

0036-8075/78/0526-0926$01.25/0

base balance, wide use of the system would improve considerably the decision-making by physicians.

Automated Medical Records

The medical record has a critical function in all phases of patient management. It is used not only by doctors but also by nurses, pharmacists, therapists, laboratory workers, and accountants, among others. Indeed, it is the link of all activities impinging on patient management and the legal documentation of all transactions concerning the patient.

PROMIS (the Problem-Oriented Medical Information System), developed by Weed (*4*) and his colleagues, is a computerized patient-management system that specifically addresses four major problems inherent in the traditional medical record. These are (i) the lack of coordination among health care providers, (ii) excessive reliance on memory by the providers of both knowledge of the patient and the medical literature, (iii) the lack of recorded rationality regarding observations and actions taken, (iv) and inadequate feedback loops for continual improvement of the practice of medicine. This system replaces paper records with the problem-oriented computer record using touch-sensitive terminal devices connected to a single minicomputer. Access to the electronic record is afforded to those involved in patient care.

Medical interactions are divided into four phases: (i) data base, (ii) problem formulation, (iii) initial and other plans, and (iv) progress notes. The data base is created by cumulating data entered for individual patients. Members of the hospital staff enter the patient's chief complaint and general physical characteristics into the system. The patient adds information on review of systems by answering yes-no and multiple-choice questions. Positive findings or responses are then reviewed by the physician, whose memory is aided by branched-logic displays, specific for each symptom or finding. The physician formulates the problem in any convenient manner. This absence of limitation allows a sign, symptom, disease complex, physiologic disorder, or abnormal x-ray finding to be identified as a problem without having to be restated in accordance with the requirement of many computer programs for a standardized language. Initial plans are entered by the physician for each identified problem. To help select among them, displays are again provided to aid in gathering data to rule out a specific cause of the problem, to gather information to assist in managing the problem, to order treatment, and so forth. The displays called up for each problem minimize the need for reliance on individual memory by providing up-to-date information from the medical literature. The developers of PROMIS believe it is this access to the medical literature coupled with instant access to specific patient data that facilitates integration and results in improved quality of care. "Progress notes" is the term applied to all subsequent notes after the initial plan whether from physicians, nurses, social workers, or others. Progress notes are entered in conjunction with data on each appropriate single identified problem. Data can be retrieved in many ways, such as on specific problems, chronologically, or in flow-sheet format. Drug orders, laboratory findings, and various aspects of the management of the patient can all be called onto the screen for immediate review.

The obvious and undeniable strength and advantage of PROMIS lies in the immediate availability of highly organized, reliable, complete records to all health care providers dealing with a patient. This alone must improve the performance of the individual health care providers as well as the coordination among the members of the health care delivery team. The organization provides a rational record relating the management action to the identified problems. The touch-sensitive input choices of displays minimizes the "computer-terminal barrier." The system's potential cost-effectiveness ratio is markedly reduced by using it effectively to support many hospital subsystems. The management summaries provide a helpful tool to review the institution's or individual's activities and practices and can contribute to a medical management-education loop. The system is not specifically directive, but in practice, 99 percent of the time the physician chooses an alternative presented by the display. Whether this is a testimony to the accuracy and completeness of the displays or to a dependence of the user is not clear.

The PROMIS user draws on a library of 36,000 displays of information (*5*). The creation and maintenance of such a volume of authoritative information covering most of medicine was and is a monumental task. The creators of PROMIS have suggested that the content of these displays could be created, maintained, audited, and validated by national groups of experts.

The Duke University Cardiovascular Information System

The purpose of diagnosis is to categorize a patient in a manner that allows the patient to be related to the physician's experience and knowledge. The physician making a diagnosis uses experience and knowledge of the disease to anticipate the course, predict the response to therapy, and influence management decisions. Deficiencies are that memory of personal experience is markedly biased by recent and dramatic events, and that the physician's knowledge is less than that potentially available.

Classification or diagnosis may suggest a homogeneous group acting in a similar manner. Often this is not the case. For example, patients with myocardial infarction may die within minutes or live for 50 years. To influence their decisions, physicians subclassify, hoping to identify one subgroup that would indeed be made up of those that live 50 years as distinct from those that die in 3 minutes. The number of subgroups the physician can carry in his head is small, however, and their characteristics inconstant. For the diagnosis of coronary artery disease, the Duke Medical Center group (*6*) estimates that at least 100 subgroups would have to be defined to achieve a reasonable likelihood that all members within each group would follow a similar course. Considering also the number of descriptors per group, the dimensionality goes well beyond the capability of the human brain. Thus, this is a problem in medicine with no solution other than by computerization.

Physicians at Duke University have developed a medical information system on the basis of their clinical experience with coronary artery disease that describes outcomes of patients with various sets of attributes (*6*). The attributes (laboratory and physical findings, history, outcomes) of a large number of patients are stored in a computer, which accomplishes the entire classification process. When attributes of a new patient are entered, the computer selects the most closely matched subgroup. The courses and outcomes of all patients previously categorized in the same subgroup are displayed. The computer's memory is wholly accurate, is unbiased by recent or dramatic events, and is enhanced by the greater number of entries derived from the entire institution's rather than one physician's experience. Therefore, the physician's management decisions can be based on far more accurate and relevant information than could be possible without the computer.

The system covers an important although relatively narrow part of medicine for which there is considerable pathophysiologic understanding. Much of the total input is hard data. The system has the advantage of continuously adding to and modifying itself. Efforts are under way to expand participation to other institutions and thus, to demonstrate its transportability. It gains authority from the involvement of large numbers of physicians participating in the cardiology care activities of Duke University. There is relatively little controversial material on which individual and consensus judgments would differ. The system is designed for application of the accumulated data to have an impact on an individual unique decision with the specifically described characteristics. Therefore, the system will improve the decision-making process of both the average and the best physician.

The Duke classification system could be used to generate homogeneous groups for whom the results of clinical trials should prove more decisive than current methods of patient selection used in such trials. Since random assignment is not the method of treatment allocation (for example, coronary bypass surgery versus medical management), differences in outcomes in subgroups have not been determined to be truly characteristic of the subgroups or characteristic of some unidentified bias in the decision-making process.

An interesting aspect of this system is that it continuously accumulates and refines data by means of which increasingly accurate estimates of risks for each intervention can be made. Thus, in part, it provides groundwork data that may be useful in the application of decision analysis to clinical decision-making.

Decision Analysis

The systems described thus far have not explicitly modeled the mechanisms of the decision-making process. Decision analysis is a formal discipline for making decisions that in many respects resembles the informal strategies of clinicians. Schwartz and co-workers (*7, 8*) present several examples of its application to clinical problems. In the differentiation of essential hypertension from functional renal artery stenosis (a surgically correctable cause of hypertension) a qualitative approach is used. The model is presented as a decision tree with several pathways, along each of which are several nodes or decision points. Each intervention is a decision node. Other "chance" nodes are not under physician control. The tree is progressively built, from signs, symptoms, laboratory tests, and diagnostic and therapeutic interventions. The tree is constructed in sufficient detail to make the representation realistic, yet is constrained to become not unmanageably large. It is acknowledged that individual physicians would differ as to the detail, but what is important is that the method forces an explicit examination and trimming of the tree by the physician. Thus, branches are pruned where both the probability and risks are low; not pruned where both are high; and assume individualized contours in intermediate combinations of probability and risk. The goal is, of course, to reach the best decision—which depends on the probabilities of various outcomes and the value judgments assigned to each outcome by the physician and the patient. But an equally important objective is to give the user far greater insight into the basis of his decisions and to force a continuing examination of his perspective and judgments.

Gorry *et al.* (*8*) have suggested that some clinical problems may be so complex that the qualitative approach is inadequate. For these situations, the procedure can be quantitated, but the calculations and representations require the aid of a computer. Such a prototypic computer program has been developed for the management of acute oliguric renal failure (*8*). This example uses disease lists, lists of attributes for each disease, and estimates of the likelihood of various potential consequences of tests and treatments. By use of the Bayesian approach, probabilities are calculated for each alternative at each decision point.

The computer program also stores large amounts of information which is available to the user. The computer's performance is, of course, consistent and separates probability estimates from value judgments. In operation, computer choices match very well those of the medical experts who created the system. It is suggested that the method may be profitably used to identify a best solution to various complex generic problems (for example, the role of gastroscopy in the management of gastric ulcer).

The methods have size limitations. Trimming the tree to retain manageability may necessitate choosing among incomplete sets of alternatives. The probabilities needed for various interventions are often not available from the literature and require use of subjective expert judgments. Gorry *et al.* (*8*) carefully acknowledged other current limitations such as a limited ability to deal with multiple diseases, inadequate dealing with the temporal changes of disease states, and the requirement of a constrained medical area.

Simulated Clinical Cognition

Pauker *et al.* (*9*) have increased the complexity of computer use in simulating clinical decision-making. Techniques of "artificial intelligence" have been merged with the decision analysis approach to develop a computer program which takes the history of the "present illness" of patients with edema.

The concept of "intelligence" is that the computer is exhibiting behavior which would be termed intelligent if such behavior were that of a person. This program is intelligent in that the computer pursues an inquiry in a manner that suggests it is trying to "understand" the illness well enough to formulate hypotheses to evaluate the clinical problem and suggest management. It does this by combining an information-gathering function with a stored information synthesis function. The latter requires no further input from the patient. The program components are (i) patient-specific data; (ii) a supervisory program that guides the computer in taking the history of the illness and generates hypotheses to be tested, accepted, or rejected; (iii) a short-term memory in which data about the patient interacts with general medical knowledge, which is, in turn, stored in (iv) an associative (long-term) memory. The long-term memory is a collection of frames each of which contains closely related facts about diseases, clinical states, or physiologic states; it consists of such information as signs, symptoms, laboratory data, and time courses, as well as rules for judging how closely a given patient's attributes match those of a frame. The frames are linked into a network by a variety of relationships (for example, "caused by" and "complicated by").

The long-term memory contains a mass of "dormant" frames of information. When some characteristics of a given frame in the long-term memory match some characteristics of the patient being explored, that frame is pulled into short-term memory (where it is then called "active") and comparisons are made between the frame's and the patient's data. As an active frame moves into short-term memory, its closely related frames (connected through a network) are pulled more closely to the short-term memory (and are thus semiactive); some of their characteristics then have greater opportunity to be compared with the characteristics of the patient. Frames that are pulled into short-term memory for comparison constitute hypotheses. Decisions are made by the program about the goodness of fit of the case to the hypothesis or on the need for more information. It measures the extent to which the frame accounts for all facts in the case, using weighted numerical scores (based on estimates of probabilities). If the hypothesis is not fully accepted, the program seeks more information by selecting questions whose answers have the greatest probability of giving useful definitive information. Cycles are repeated until the program has completed its diagnostic process.

The investigators believe that this procedure represents an important advance because it retrieves and applies knowledge when required, thus freeing the programmer from the impossible task of specifying all contingencies in advance, as would be required in branching flowchart approaches. The goal-directed character of the computer supervisory program allows (i) selection of the pertinent medical and real-world knowledge from the computer's memory and (ii) dynamic assembly of small problem-solving techniques to guide acquisiton of additional information.

The investigators' assessment of their efforts identify certain persistent limitations of the system. The most important are the limited ability to deal with a wide range of clinical problems rather than with constrained areas and the lack of capacity to deal with coexisting diseases or to fully use the temporal and changing aspects of the patient's history and course. Better methods are needed to aid the computer to focus rapidly on a narrow range of concern, a function now performed very well by the experienced physician.

INTERNIST: A Problem Solver

INTERNIST is a computer-based diagnostic consultation system for problems in internal medicine (*10, 11*).

If a full range of diagnostic categories had to be specified in advance, the number required to arbitrarily classify patients has been estimated at the order of 10^{40} (*11*). This is not because a large number of disease entities are known to a clinician (estimated from 2000 to 10,000), but rather because of the presentation of concurrent clinical problems (a dozen or more is not rare). Given what may appear to be an indiscriminate collection of data (signs, symptoms, laboratory results, for example), the clinician's first job is to decide what problems are to be dealt with. J. D. Meyers and his associates at the University of Pittsburgh have concentrated on a model of synthetic reasoning that simulates those aspects of an expert clinician's behavior concerned with formulating composite problem hypotheses.

The system is based on assigning rough estimates of likelihood of the association of a disease given a finding and a similar estimate of the likelihood of the finding given a disease. Thus, each disease entity has an associated list of manifestations known to occur in that disease and an estimate (on a scale of 1 to 5) of the frequency of occurrence of each in that disease. Similarly, for each manifestation there is a list of diseases in which it is known to occur and a weighting factor (0 to 5) for each. The weighting factor is not a simple estimate of frequency but rather an estimate of strength of association or likelihood that the manifestation and disease are causally related. This weight is called the "evoking strength." There is also recorded a hierarchy of disease categories, which has at the top level primarily organ systems (such as kidney disease, lung disease). Each of these may be subdivided any number of times until a specific disease entity is reached (chronic glomerulonephritis, carcinoma of the lung). Other accessible information in the data base includes causal, temporal, and other associations by which various disease entities are interrelated. The estimates developed for the knowledge base are based on general medical knowledge, a review of the literature, discussion and consultation with subject-matter experts, clinical case experience, trial use of the system, and a variety of other sources, all of which are funneled through and ultimately decided on by Meyers. This is a continuing process.

In the scoring process, the evoking strength is particularly useful in reducing the dimensions of concern, thus acting as a "constrictor." The evoking strength and the importance of manifestations explained by a disease being tested are counted in its favor; frequency weights are counted against those diseases in which the manifestations are expected but not found in the problem case. The INTERNIST system uses constrictor concepts to delineate top-level structures of complex problems and considers within subproblems only those findings that are relevant. Multiple problems are addressed simultaneously, with the likelihood of each disease under consideration being indicated at each step. The system presents a running record of those findings not explained by the hypotheses being considered. To pursue the diagnosis, the program identifies additional manifestations (symptoms, findings, laboratory tests) about which it would like additional information in order to reach a more definitive conclusion. In so doing, it weighs the risks, hazards, discomforts, and other adverse aspects of gaining such information. It asks first for simple data that can be obtained without much cost. The program can reach conclusions about simple problems (a single disease) or complex mixes of diseases and their associated disease states.

It is estimated that the system now covers about two-thirds of internal medicine. When the test patient's problems are within the system, informal testing of the system's performance has been remarkable. Obviously, it cannot diagnose a disease not yet entered into its knowledge bank. The system mimicks the diagnostic behavior of the excellent clinician. It can partition manifestations into those associated with one or more disease states and leave unsettled and unexplained some manifestations (as is frequently done when the judgments are made by the expert clinician).

The system covers a broad area of medicine. It can deal effectively with areas where the pathophysiology is little understood and where much of the data is soft. New information is constantly added. It effectively uses a data base far greater than that retainable by any individual physician. The system retains the episodic peaks of excellence reached by the creator and his colleagues at the time the information is entered and does not suffer from what Meyers calls the "fadeaway" of information with the passage of time. In his view, the system can therefore do things he himself cannot.

The authority of the knowledge base derives from the input of a small group of physicians, and ultimately from the broad experience, competence, and synthesizing ability of one physician. The incredible complexity and ever-increasing volume of new medical information suggests that the ultimate development and maintenance of a complete knowledge base will involve developing some form of consensus by multiple experts in their respective subspecialty areas.

A Medical Information Bank

Each of the medical information systems described so far intrudes into and participates directly in the interactions between the physician, the patient, patient data, and general medical knowledge as decisions are made. An alternate approach is to develop methods for making readily available the current state of that medical knowledge relevant to the specific decisions being made. Such an approach focuses on remedying existing deficits in the transfer of biomedical information and is directed toward extending and augmenting the physician's memory and capacity to synthesize the vast and ever-increasing medical literature, but it does not intervene directly in the decision-making process. Reliance is placed on the physician (or other health care provider) to use appropriately the knowledge provided in making decisions. An effort of this kind is not a substitute for the medical information systems described. Rather, it is a prototype method for developing authoritative knowledge bases.

Within the Lister Hill National Center for Biomedical Communications (the research and development arm of the National Library of Medicine) a program is aimed at developing, demonstrating, and evaluating a prototypic, computerized information transfer system. As a comprehensive bank of information, the system will (i) contain substantive answers to questions posed by practitioners, (ii) provide answers that are current and are the consensus of a group of experts, (iii) be immediately responsive to inquiries, and (iv) provide data supporting the answers as well as citations to primary publications for more detailed study if desired. The diseases "viral hepatitis" have been selected to serve as the initial test model.

A data base suitable for automated search and retrieval techniques has been constructed. Knowledge pertaining to aspects of viral hepatitis important to the practitioner has been synthesized from the information contained in several outstanding review articles previously published by hepatitis experts (*12*). Relevant information has been selected, placed in a highly organized hierarchical arrangement to permit easy retrieval, and encoded into a minicomputer. The data base is arranged by topic headings. For each heading there is an accompanying heading statement that synthesizes the state of knowledge about the subject. Data elements—paragraphs taken from previously published sources—support each heading and heading statement. Citations included within the data-element paragraphs are to the primary publications cited by the experts in their review articles to support their conclusions or general statements.

The synthesized statements offered by the information bank reflect a consensus of nationally recognized experts on the state of knowledge of the subject, and may convey substantive unanimous agreement, support for two or more mutually exclusive views, or indicate that information in a given area is simply lacking. By consensus, the same group of experts will maintain currency of the data bank by monitoring selected newly published material and revising as appropriate. A "computer-conferencing" network, linking the geographically dispersed experts with one another and with the staff of the National Library of Medicine, will serve as the principal medium of communication to facilitate monthly updating of the data base.

Planned access to the information bank may be available directly through a computer terminal or indirectly through a trained intermediary who can be reached on a toll-free telephone. Users may also have telephone access to programmed questions and answers or receive computer-generated printed material in hard copy or on microfiche. Other derived multimedia products would also be available.

Discussion

The medical information systems described demonstrate important advances in the use of computers to help the physician make decisions. Computers clearly can function successfully as extenders of the physician's memory for information generally and for individual patient data, and they can greatly augment the information resources available to support decision-making. The approaches are in some ways complementary, in other ways overlapping; in no way can the sum of them be considered to address fully and successfully the recognized problems in the arenas of diagnostic, prognostic, and treatment decisions.

Four of the systems perform within narrow ranges of disease but in considerable detail. They are the electrolyte and acid-base disorders consultation system, the Duke coronary artery system, decision analysis, and the computer simulation of clinical cognition. In contrast, PROMIS and INTERNIST perform in less detail, but across broad areas of disease.

The electrolyte and acid-base disorders computer consultation system and the Duke coronary artery system cover small substantive areas of medicine for which hard data provide much of the critical input and about which considerable pathophysiologic understanding exists. Because of their highly constrained nature, they can operate on the basis of a relatively small number of definitive diagnoses. Although this approach can be extended to some other constrained arenas of medicine, there are many substantive areas in which these criteria are not met. Such a system can be useful only after the physician recognizes that the patient's problem falls within its scope. These approaches cannot deal simultaneously with arrays of multiple disease categories as can the expert clinician, nor do they effectively deal with the temporal relationships of additional data.

Systems using decision analysis and computer simulation of clinical cognition share most of the characteristics described above. Although applied so far only in constrained areas, these approaches could be extended to all medical areas. They alone offer very explicit consideration of value judgments assigned to outcomes by both physicians and patients; all other systems rely upon implicit integration of both the physician's and the patient's value judgments in the decision.

Both PROMIS and INTERNIST cover broad areas of medicine, including those for which the pathophysiology is little understood and much of the data soft. As the scope increases, diagnostic and therapeutic directiveness decreases. The breadth of coverage requires developing and maintaining an authoritative data base, which can best be accomplished through the collaboration of a large number of subject specialists. PROMIS deals with multiple disorders largely by treating them separately and sequentially. INTERNIST mimics an expert clinician and is the only one of these systems that handles symptoms and findings of multiple diseases simultaneously.

The medical information bank approach is a prototypic attempt to make readily available an authoritative statement of current medical knowledge relevant to specific decisions being made. The information can be used directly, or it can complement other computer approaches to medical decision-making.

Computers perform as instructed in support of the various systems described. The acid-base program assumes that each of its diagnoses is exclusive; it deals with management on this unitary basis by a completely preprogrammed management flow-sheet display. The computer problem is reduced to choosing the appropriate display. The computer problem for PROMIS is qualitatively the same, but now the choice is from a library of 36,000 displays. In the coronary artery system the displays are being revised constantly as new experience is gained. These revisions, made by the computer, provide directly authoritative information necessary to make decisions on subclassification (diagnosis), prognosis, and managment for individual patients. Computer simulation of clinical cognition has introduced an important "synthesis" concept. Here the final displays do not exist in a preformed set but must be built up by appropriate combination of small modules. The computer problem is to select and apply knowledge from its stores when it is required, thus freeing the programmer from the impossible task of prior specification of all possible contingencies. This concept of synthesis is also employed in INTERNIST which combines it with elements of probabilistic computations and pathophysiologic flow charts to build and select displays. INTERNIST keeps track of and displays both explained and unexplained findings. The hepatitis data base uses the computer for simplifying text editing and updating and for convenience of storage, retrieval, and dissemination.

In this brief review of examples of some of the classes of computer systems being developed to support diagnosis, prognosis, and therapy, there is evident a clear evolution of both an increasing sophistication of systems and a progressive recognition of the complexity of the problems. Also changing is the man-machine relationship. What may have been considered earlier as an adversary relationship is evolving through greater recognition and respect for the unique capabilities of each into a synergistic collaboration. Whatever the limitations of existing systems, the data justify an optimistic view of the future of this collaboration in medicine.

References and Notes

1. J. R. Senior, *Toward the Measurement of Competence in Medicine* (Rittenhouse, Philadelphia, in press).
2. H. L. Bleich, *J. Clin. Invest.* **48**, 1689 (1969); *Am. J. Med.* **53**, 285 (1972).
3. ——, *N. Engl. J. Med.* **284**, 141 (1971).
4. L. L. Weed, *Your Health Care and How to Manage It* (Essex, Essex Junction, Vt., 1976).
5. *NCHSR Research Digest Series: Automation of the Problem-Oriented Medical Record* [Publ. (HRA) 77-3177, Department of Health, Education, and Welfare, Washington, D.C., 1977].
6. R. A. Rosati, A. G. Wallace, E. A. Stead, Jr., *Arch. Intern. Med.* **131**, 285 (1973).
7. W. B. Schwartz, A. G. Gorry, J. P. Kassirer, A. Essig, *Am. J. Med.* **55**, 459 (1973).
8. A. G. Gorry, J. P. Kassirer, A. Essig, W. B. Schwartz, *ibid.*, p. 473.
9. S. G. Pauker, A. G. Gorry, J. P. Kassirer, W. B. Schwartz, *ibid.* **60**, 981 (1976).
10. H. E. Pople, J. D. Meyers, R. A. Miller, in *Proceedings of the 4th International Joint Conference on Artificial Intelligence (Tablisi, USSR, 1975)*, pp. 848–855.
11. H. E. Pople, in *Proceedings of the 5th International Joint Conference on Artificial Intelligence (Cambridge, Mass., 1977)*, pp. 1030–1037.
12. J. W. Mosley and J. Galambos, in *Diseases of the Liver*, L. Schiff, Ed. (Lippincott, Philadelphia, ed. 4, 1975), pp. 500–593; National Academy of Sciences Symposium on Viral Hepatitis, *Am. J. Med. Sci.* **270** (Nos. 1 and 2) (1975).

The Road from Research to New Diagnosis and Therapy

Julius H. Comroe, Jr.

My career goal was teaching—not research. When I finished my internship, I became an instructor in pharmacology at the University of Pennsylvania because it was the best teaching department in the School of Medicine and I wanted to be part of it. (It was also the best research department, but that was a secondary factor in my decision.)

One learns a great deal by teaching, especially if one has bright, inquisitive, uninhibited students. The first thing that I learned (and I learned it the first day) was that I did not know the answer to many of their questions. The second thing I learned was that, for most of their questions, no one else had good answers either. In short, the areas of ignorance were far greater than instances of solid, real knowledge. That was in the mid-1930's. We could determine the specific type of pneumonia bugs in a patient, but we could not treat the patient because we had no sulfas or penicillin. We had x-rays that could detect shadows in lungs, but except for the vital capacity test, we had no tests of the function of the lungs. We had iron lungs (even built for two) for polio patients with respiratory paralysis, but we had no vaccine to prevent the disease in the first place. We could diagnose pulmonary tuberculosis, but we had no way to cure tuberculous patients because we had no streptomycin or para-aminosalicylic acid or isoniazid. We thought we were helping (or possibly curing) patients with tuberculosis by insisting on two plus years of bed rest in sanitoriums and by pneumothorax, but we know now that we really were not. We knew when a patient was not breathing and needed resuscitation, and the Red Cross had taught everyone in the country how to use the Schafer method of prone-pressure artificial respiration—until someone made actual measurements of the volume of "good air in" and "bad air out" and found it inadequate to sustain life; only then were adequate methods devised.

We could measure high blood pressure in the systemic circulation and diagnose hypertension there, but we could not measure pulmonary arterial blood pressure and diagnose pulmonary hypertension because Cournand and Richards had not yet done their basic studies of cardiac catheterization. We had no methods at all in the mid-1930's for resuscitation of the stopped heart and no artificial pacemaker to make a too-slow heart beat at a normal rate. We could tell, by looking, when newborn babies could not breathe properly and were blue, but we did not know why and most of them died; now research has shown us what causes the respiratory distress syndrome, and knowing that, we can treat it effectively and most of these babies now live.

My father, practicing internal medicine in the mid-1930's, had a few drugs (digitalis, insulin, arsphenamine, vaccines) that improved, cured, or prevented disease, and a few more (aspirin, morphine, barbital) that relieved symptoms of disease. Mainly he provided his patients with hope, encouragement, relief of suffering, and laxatives to keep their bowels open, and he recommended excision of foci of infection (mainly tonsils) and plenty of fruit juices (he had no faith in chicken soup). I, teaching in the mid-1930's, had few honest answers to

The author is Morris Herzstein Professor of Biology and Professor of Physiology in the Cardiovascular Research Institute, University of California, San Francisco 94143.

0036-8075/78/0526-0931$01.50/0