AI *and the* **Changing Face of Health Care**

Erika Rogers

The last few years have seen growing optimism that the world of medicine will finally be truly online. Many believe that computerization is the answer to some of the major challenges faced by the health care community. Also, because of increased networking capabilities, effective new solutions to old problems appear to be emerging. Furthermore, both administrators and practitioners are becoming convinced that technology will change the face of health care, balancing improved patient care with cost-effective management procedures.

AI's role in provoking and supporting these changes is of particular interest because, despite AI's long history of research in the medical domain, relatively few AI systems are in, or are in reach of, clinical use. The first part of this article, therefore, examines the relationship between health care and AI today. However, the future promises a greater role for AI in Medicine (AIM). To support this prognosis, the remainder of the article addresses these questions:

- What are the new directions for health care changes?
- How do these changes affect AI work in this area?
- What can AI contribute toward realization of these changes?

AI and health care today

The application of knowledge-based techniques in the medical domain is a familiar and time-honored theme. Since the 1970s, medical projects such as Mycin, Casnet, Internist, PIP, Digitalis Therapy Advisor, IRIS, and Expert have been in the spotlight for two major reasons (for more background on AI in Medicine, see the related sidebar). One was that the medical domain offered a rich environment for research and development of virtually every aspect of knowledge-based theory. This found expression in two types of interaction:¹

- Technology-driven research, where medicine provides AI with a good set of problems that are used to develop techniques leading to more general-purpose technologies. These can then be applied to other domains (for example, diagnosis of computer circuits).
- Inward-looking research, which addresses technical solutions that, although neither directly AI-based nor of immediate concern to clinicians, are needed to support the long-term goals of both enterprises. One such problem is the design and implementation of integrated electronic patient records.

In both of these types of interactions, AIM has proven exceptionally successful. Medicine has provided AI researchers with a fertile application domain that is still, decades later, nowhere near exhaustion.

The second reason these projects received

such attention was that AI was expected to be of reciprocal benefit to the field of medicine. Optimistic predictions claimed that AI technology would contribute greatly to medical practice by assisting, augmenting, or even replacing the decision-making process. Such claims have tended to color the assessment of AI's overall value to the medical community. After so many years of effort and research, many people, including AI researchers themselves, are asking whether AI has made any identifiable contributions to clinical medicine. Its impact on health care appears to be relatively minute, with a few individual systems in clinical use scattered among hospitals and practices throughout the world. Although this part of the vision was (and still is) grand, the time frame for achieving it has needed some revision. Edward H. Shortliffe succinctly alluded to this in 1993 when he referred to "the adolescence of AI in medicine."²

Inside problems. Although numerous causes for this apparent lack of progress have been proposed, the roots of the problems probably lie in the separate challenges faced by AIM researchers and medical professionals. AIM researchers attempt to solve difficult computational problems. Examples include the acquisition and representation of a practitioner's knowledge and skills, integration of episodic as well as longitudinal and historical patient data, and communication with real-time monitors and devices.

Medical professionals, on the other hand, not only must acquire and process knowledge, but are also subject to the human factor. The impact of diagnosis and treatment on patients' lives, the threat of professional liability, and the pressures of administrative and financial constraints are just a few of the additional issues that define the viewpoint of a medical practitioner.

These differences have sometimes led to conflicting or contradictory goals in the partnership between AIM researchers and their medical collaborators. AI has also frequently been hampered by its own concerns and limitations. Handling the huge corpus of well-established medical knowledge, adapting to the dynamic and uncertain nature of medical practice, and discovering and modeling how physicians successfully perform diagnosis at all, are problems that are not necessarily germane to the medical community.

Outside problems. In addition to these "self-imposed" challenges, external problems have

A brief background of AI in medicine

For an introduction to some of the classical AI systems in medicine, volume two of *The Handbook of Artificial Intelligence*¹ is a good place to start. A more recent publication is Randolph A. Miller's threaded bibliography and commentary.² This 1994 article presents an excellent overview of medical-diagnostic decision-support systems and their principles, as well as a chronological bibliography of 40 years of research in this area. Mario Stefanelli has written an overview of some of the AIM work in Europe.³

The journal *Artificial Intelligence in Medicine* is devoted to research issues and results in this area. It frequently has special issues devoted to topics such as case-based reasoning, temporal reasoning, and ontological issues. Proceedings from conferences such as AIME (Artificial Intelligence in Medicine Europe)⁴ present the most recent results in a broad range of topics such as

- protocols and guidelines,
- knowledge acquisition and learning,
- decision-support theories,
- diagnostic problem solving,
- probabilistic models and fuzzy logic,
- temporal reasoning and planning,
- natural language and terminology,
- image and signal processing, and
- hybrid and cooperative systems.

For a more detailed set of resources, access <http://www.csc.calpoly.edu/~erogers/AIMED/aimed-resource.html>, which contains links to major topics in AIM and contacts for newsgroups, mailing lists, and other relevant information.

References

1. A. Barr and E.A. Feigenbaum, eds., *The Handbook of Artificial Intelligence, Vol. II*, William Kaufmann, Los Altos, Calif., 1982.
2. R.A. Miller, "Medical Diagnostic Decision Support Systems: Past, Present, and Future," *J. Am. Medical Informatics Assoc.*, Vol. 1, No. 1, Jan./Feb. 1994, pp. 8-27.
3. M. Stefanelli, "European Research Efforts in Medical Knowledge-Based Systems," *Artificial Intelligence in Medicine*, Vol. 5, 1993, pp. 107-24.
4. E. Keravnou et al., eds., *Artificial Intelligence in Medicine, Lecture Notes in Artificial Intelligence 1211*, Springer-Verlag, New York, 1997.

affected AI's performance and use in clinical settings. Two of the most commonly cited are lack of user acceptance and lack of infrastructure. The first is really a two-way problem (both parties sometimes accusing each other of "arrogance"). Physicians and other medical practitioners are concerned about using automated systems that try to take over the decision-making process without substantiating the suggested decisions, and that are frequently based on very limited knowledge. The systems only work on a small subset of the patient cases and often have user interfaces that interfere with the workflow of processing and treating patients. This concern is completely justified when we consider the seriousness of the decisions to be made. Lives are frequently at stake, and practitioners—not the AI systems—are still the ones held accountable for those decisions. The other side of the acceptance coin is that some of this attitude is based on an outmoded view of AIM—the old "Greek oracle" system has long been replaced by a more user-centered approach, and recent systems are much more usable and useful.

However, persistent negative stereotypes of "expert systems" might also be related to prejudices against computer-based systems in general. Many problems that cause such resistance stem from a basic lack of infrastructure. For example, information about patients has typically been kept in multiple databases, multiple formats (from digital to handwritten), and multiple departments. This has led to a kind of chicken-and-egg problem: To do knowledge-based work, AI researchers typically have had to construct their own version of an electronic patient record first, try to populate it with actual patient data (often by hand), and then start to operate on it in a more interesting fashion. Unfortunately, this often involves tradeoffs. In some cases, researchers limit the amount of actual patient data and test the AI algorithms on a reduced or artificially constructed set of cases. On the other hand, tackling the more general patient-record problems can lead away from the original knowledge-based goals. Grants run out, researchers change focus, people lose interest, and just when the system could start demonstrating deeper intelligence (perhaps 10 years later), the project is over. Shortliffe has described this problem most eloquently:

Resistance to system use has occurred despite the inherent merit of the methods that AIM researchers have developed. Many of the problems are, instead, a reflection of the disarray of our health care system, and, more importantly, its failure to build the local, regional, national

and international infrastructure for biomedical computing and communications that will be required before computer-based decision-support tools can become routine elements in the clinical setting.

...Our limited success with dissemination of knowledge-based systems in medicine is, in my view, due more to this failure of integration than it is to any other basic problem with the AI technologies that have been developed.²

The shape of things to come. The list of challenges continues, but before we become mired in the difficulties, let's look at what is happening right now in health care. Words such as computers, telecommunications, databasing, outcomes management, and expert systems are suddenly on everyone's lips—not just as impressive jargon, but as concrete aspects of the immediate medical future.³ These aspects will have major implications for the AIM community.

The changing face of health care

The state of health care is on everyone's mind, and rightly so—it touches all aspects of our lives. When we or our loved ones need medical assistance, we want it to be fast, effective, and preferably at a manageable cost. Therefore, increasing quality while reducing costs is a major concern. Nowhere is that more evident than in the United States, where health bills consumed 14% of the gross national product in 1995 (approximately \$1 trillion).⁴

Extreme views suggest that practitioners should accept the inevitability of managed care and its attendant "advantages":

for all the moaning and groaning about the depersonalization of medicine, the deterioration of the doctor-patient relationship, and the attack of the killer HMOs, physicians who want to practice good medicine in the future, and to give good advice to patients, will have to be integrated into health care organizations, and will have to accept computerized decision making.⁵

Even moderate views agree that major changes can and must take place.

A brief look at four specific areas of change will provide insight on where health care is going and how AIM might foster these advances. The first two, *infrastructure* and *medical imaging*, are primarily technology-driven, while the second two, *patient-centered health care* and *evidence-based medicine*, are driven from within the medical community itself.

Infrastructure. Regardless of the political and financial structures of the health care establishment, various experts believe that computer and networking technology is the key to the revolution that is forging the future:

I would posit that between now and 2020, we have a grand challenge. ... we can attend to that challenge only by capitalizing upon information and communication technology.⁶

Market growth will depend on the ability of the health care and information systems industries to solve the "many to many" problem: integrating many pieces and types of information, in many formats, on many platforms, from many stakeholder environments, for use in many geographic locations.⁷

Advances in health-oriented telecommunications, medical imaging, massive databasing, memory miniaturization, satellite technology and other information systems lay the groundwork for fundamental changes in the organization of health care.³

Many are looking to the National Information Infrastructure (NII) with its "widely accessible and interoperable communications networks"⁷ to solve the health care dilemma of cost versus quality by providing access to digital libraries, information databases, and decision-support services. The Internet is already having an impact, as massive amounts of information are becoming available anywhere, anytime, to anyone. Although this democratization of information dissemination is laudable, major concerns have arisen regarding security, privacy, and quality control (for example, how can an Internet user distinguish between useful, valid medical information and something posted by an unqualified amateur?). Discussions of these issues must take into account the viewpoints of the patient, the public, medical practitioners, and health care administrators. The establishment of a reliable computational foundation, both at individual medical centers and more globally, will set the stage for additional fundamental changes in the way medicine is practiced.

This will also provide a new starting point for the development of AIM systems. Algorithms and theories can be implemented and tested without the previous qualitative and quantitative limitations on patient-data access. General medical knowledge will be more easily integrated, and new research questions will be explored in the development of intelligent software agents to mine this information and data for automated knowledge acquisition.

Medical imaging. The ability to see inside the body through imaging technology has

long been a fundamental part of medical diagnosis and is one key to a patient's journey through diagnosis and treatment. Physicians and radiologists use images to identify and label abnormal processes in the patient's body; surgeons use images to plan the paths of their instruments and to develop strategies for dealing with the physical location of the abnormality; nurses use images to predict pre- and postoperative signs and behaviors that will guide their care. Several basic imaging techniques, such as computed tomography (CT) and magnetic resonance imaging (MRI), produce digital images. However, hospitals and practices that do not have digital picture archiving and communication systems (PACS) typically generate images in the radiology department and make hard-copy films, which accompany the patient's chart as it moves through the care process.

With the introduction of high-bandwidth communications, image generation and management is taking on a new perspective. Multispecialty teams will be able to easily examine and diagnose images, as well as exchange reports and data on these patients.

One of the earliest areas to take advantage of NII-style networking has been *teleradiology*: the transmission of images between distant sites. Such systems enable primary-care physicians at hospitals and clinics in rural or remote areas to interact and consult with experts at major centers. Typically, digital copies of film (or sometimes actual digital images) are transmitted, and then an expert is paged to examine the images and to discuss the case by telephone. This technology has led to major savings by reducing unnecessary patient travel (decisions can be made not to transport a patient to a major center) and by expediting decisions calling for urgent transport when the case is serious.

An interesting side effect of this work is user acceptance. For many years, there has been an open debate about whether the reduced quality of digital images compared to film images has a negative effect on diagnostic ability. Yet images from teleradiology systems, which typically have substantially lower spatial and contrast resolution, are nevertheless accepted and used to make diagnostic decisions. This suggests that medical professionals do embrace new technology when the benefits to patient care outweigh the apparent disadvantages.

The increased availability and accessibility of image data in electronic form provide extensive opportunities for AIM research to make

major contributions to both in-house and remote radiological practice. Techniques from computer vision as well as knowledge-based image-retrieval, processing, and enhancement algorithms can be combined with intelligent display and diagnosis capabilities. These kinds of systems will greatly enhance the decision-making support available to physicians, radiologists, surgeons, and nurses, and will ease the bottleneck of current image-handling practices.

Patient-centered and community-based health care. Another major trend in health care that the NII will promote is toward patient-centered and community-based medical practice. This will require more patient involvement in the therapeutic decision-making process:

The changes under way in the health care arena are compelling patients to assume more responsibility for their own health care and the care of loved ones. Providers are encouraging this process by promoting wellness programs and delivering more health-related information to their patients.⁷

To make informed decisions about preferences, concerns, and values, the layperson must be able to obtain relevant medical information that is understandable, from sources that are easily accessible.

The focus will move out of acute-care hospitals, and back to clinics, to doctor's offices, and even into schools, workplaces and the home. ... the focus will also change from intervening in the acute phase of the disease, and toward early screening, detection, and treatment.³

Courses such as Harvard's "Cybermedicine—The Computer as a Patient's Assistant" are providing training for this community-based effort:

The premise is that the largest, yet least used, health care resource worldwide is the patient or prospective patient, and that interactive computers that enlighten patients will thereby improve the quality of their care. A corollary of this premise is that the cost of medical care can be reduced substantially if the computer increases the patient's ability to participate more fully in the process of care.⁸

And what seemed, until recently, futuristic visions are now close to realization:

We could save money in big buckets if we gave people an easy way to grab good information about their own health at home. ... Picture the kind of TV/telephone/computer information appliance that people widely expect will be a big part of the home in the future. Give it a home version of a medical expert system, fitted with a highly interactive graphic interface. ... Based on what you tell it, the system can triage the cases you can take care of yourself from the ones that require a doctor's care, and the ones that require instant attention.⁹

However, although the underlying networking and hardware capabilities will soon be in place (at least in first-world countries), a great need exists for the design and development of the software programs that will make this a reality. Medical knowledge must be captured and presented in a totally different way from that used by the medical profession. These systems must bridge the communication gap between professional and lay terminology, as well as provide useable and intelligent interfaces for a large variety of people. Such issues as speech, natural-language processing, and adaptive interfaces that are customized to the user are of particular interest to the AIM community.

Evidence-based medicine. This trend is receiving international attention as a way of transforming the nature of clinical practice.¹ David L. Sackett and his colleagues have defined it as

the conscientious, explicit and judicious use of current best evidence in making decisions about the care of individual patients. The practice of evidence-based medicine means integrating individual clinical expertise with the best available external clinical evidence from systematic research.¹⁰

Although controversy surrounds this approach in the medical community, it has aspects that appeal to AI researchers. For example, it encourages "systematic attempts to record observations in a reproducible and unbiased fashion."¹¹ Such a process will supplement the physician's clinical experience. It will also provide a more reliable method of developing clinical instincts and knowledge about patient prognosis, the value of diagnostic tests, and the efficacy of treatment. This approach also includes the development of practice guidelines based on a rigorous methodological review of the external clinical literature. Several medical schools now include training in evidence-based medicine, which requires new types of analytical skills not traditionally taught, including

precisely defining a patient problem, and what information is required to resolve the problem; conducting an efficient search of the literature; selecting the best of the relevant studies, and applying rules of evidence to determine their validity; being able to present to colleagues in a succinct fashion the content of the article, and its strengths and weaknesses; and extracting the clinical message, and applying it to the patient problem.¹¹

This area promises two main benefits to AI researchers. First, it lays the foundation for the acquisition of medical knowledge in

a more systematic way by the practitioners themselves. (Anyone who has worked in this domain knows the challenge of acquiring even "well-documented" knowledge, which is frequently in free-form, hand-written physician's notes or in proprietary databases with limited access and limited information. Acquiring the deep knowledge that a physician uses constantly in weighing observations and evidence to make diagnostic decisions poses an even greater challenge. These concepts have traditionally been extremely difficult for many practitioners to articulate and for AI researchers to then reformulate into a computational form.) The second benefit is somewhat more subtle: training physicians to think and practice in this more methodological manner might also make them more amenable to working with software systems that reason in the same way.

The changing face of AIM

Particularly interesting and, in some ways, surprising is that in all these areas, every expression of the "vision" explicitly presumes that AI technology will play a major role. One prediction speaks of "unlimited low-cost computing power providing physicians with Artificial Intelligence support systems that they will be able to download from easily accessible software networks."¹² Another expects that "individuals will have access to community aggregate health data" and that "patients will be able to decide whether or not to contact a physician through AI programs, informed by outcomes data, and accessed through their home entertainment modules."¹² Don E. Detmer has emphasized the development and management of knowledge bases as an important part of meeting the next century's grand challenges.⁶ Visionary articles on health care trends also include knowledge-based or "expert" systems as a crucial part of the wave of the future:

A researcher will be able to ask "what if" questions, and answer them by dipping into databases with hundreds of millions of records, and aggregating all cases that fit. ... Combine such databases with expert systems and you get expert systems that can learn. ... Tie the health care system into a vast data structure, and you get a system that can learn how to do it better, faster, cheaper and easier.⁹

Medical knowledge is expanding faster than any human can learn it. Computer programs called expert systems help physicians and other health practitioners move much more rapidly and effectively through the decisions of diagnosis

and therapy, isolating rare diseases, differentiating between similar syndromes, and discovering the latest research on the most effective therapies. Their widespread use is likely to significantly change the role of a doctor away from knowing facts and toward the more human elements of the craft, such as making difficult judgments, and helping patients change their behavior.³

New challenges. The new infrastructure will support networks of large knowledge bases of patients, together with "inferential tools in the clinician's system [that] make recommendations on testing and diagnosis, guided by cost-benefit considerations, balancing the cost of testing with the benefits of early treatment of possible disorders."⁷ These tools' predicted capabilities include

- helping to assess a patient's current or future condition, given the clinical background and symptoms, and to select the best therapy under uncertain conditions;
- searching for and tracking common and important errors in diagnosis and therapy, to learn patterns of error from a large body of clinical interactions;
- detecting costly testing and treatment strategies and redirecting physicians to more cost-effective approaches; and
- appraising health risks and monitoring health care performance: looking for missing information, identifying clinical protocols that match the patient's, and providing computer agents to conduct such chores as locating the most cost-effective MRI or CT scan in a neighborhood or preparing a cost-benefit analysis on proposed tests or treatments.⁷

AI researchers can use such predictions as a roadmap to determine what the medical community considers interesting problems.

New opportunities. All these changes in health care bode well for AIM research. Shortliffe has said that

the greatest hope for effective systems will be realized when the infrastructure for introducing computational tools in medicine has been put in place by visionary leaders who understand the importance of networking, integration, shared access to patient data bases, and the use of standards for data exchange, communications, and knowledge sharing.²

Such visionary thinking is already happening in the health care community, although change will not be instantaneous, by any means. In the meantime, the AIM

community must also reevaluate its priorities. New approaches to old problems might more likely gain acceptance. Advances in medical technology are opening new and broader areas of research. Most important, AIM research must take into account a major shift in attitude: The medical community not only strongly desires knowledge-based solutions, but also expects and demands them.

AS THE ARTICLES IN THIS SPECIAL issue (see the related sidebar) demonstrate, AI can and will make a difference: first, to clinical medicine, by placing high-powered assistive tools in the hands of the medical experts; and second, to general health care practice, with its emphasis on preventive care and patient education and empowerment.

Contributions from the AI community are beginning to receive recognition. These contributions range from the development of standard ontologies to the provision of expert assistance in the form of decision-making rules, workflow support, just-in-time training, libraries of reusable knowledge, and so on. However, the large scale of these problems means that they can no longer be addressed by isolated groups of specialists. Network and database technologies must be combined with knowledge-based techniques to enable the health care industry to manage patient information not only during clinical episodes, but also over the lifetime of patients and their families.

As both the health care industry and technology researchers address these problems, AI and knowledge-based applications should come into their own, providing powerful software systems for such tasks as diagnostic decision-making, patient monitoring, management planning, and outcomes prediction, to name just a few. To paraphrase Enrico W. Coiera, the question in the new millennium will not be, "What has AIM achieved lately?" but rather, "How did we ever live without it?"¹ ■

References

1. E.W. Coiera, "Artificial Intelligence in Medicine: The Challenges Ahead," *J. Am. Medical Informatics Assoc.*, Vol. 3, No. 6, Nov./Dec. 1996, pp. 363-366.

In this special issue

The five articles and two applications pieces in this special issue present a snapshot of the state of AI in Medicine. They illustrate some of the important issues, while offering a view of AIM's future.

In "Integrating a Knowledge-Based System for Parenteral Nutrition of Neonates into a Clinical Intranet" (pp. 65-69), Werner Horn, Christian Popow, Silvia Miksch, and Andreas Seyfang address the practical problem of calculating and planning nutrition solutions for sick and premature babies. Their system works on an intranet of hospital workstations with an html interface, and has achieved a large measure of acceptance. One major challenge they faced was to integrate their AIM application with the hospital's regular patient-data-management system.

On a different scale is the project discussed by Jan Eric Larsson and Barbara Hayes-Roth in "Guardian: An Intelligent Autonomous Agent for Medical Monitoring and Diagnosis" (pp. 58-64). Their system addresses many problems in real-time intelligent monitoring and control of intensive care for open-heart surgery patients. To test Guardian, Larsson and Hayes-Roth used both its internal simulator and an external simulator normally used in a training studio for human medical personnel. Their work typifies the enormous effort that must go into an AI system that addresses large-scale problems in complex medical environments.

In "Neural Network Learning for Intelligent Patient Image Retrievals" (pp. 49-57), Olivia R. Liu Sheng, Chih-Ping Wei, and Paul Jen-Hwa Hu address the problem of retrieving relevant patient images and discuss the application of back-propagation neural network techniques to automatically learn how radiologists make these decisions. Their work exemplifies the promising contribution that AIM can make to medical-image management.

On the other hand, in "Knowledge Architectures for Patient Access to Breast-Cancer Information" (pp. 26-31), Colleen Crangle, Robert Carlson, Lawrence Fagan, Mark Erlbaum, David Sherertz, and Lauren Langford present a patient-centered approach. Their system provides patients with online access to FAQ-based information about breast cancer and was adapted from a knowledge-based information system originally designed for clinicians. They are testing it in a community project.

"TraumaTIQ: Online Decision Support for Trauma Management" (pp. 32-39), by Abigail S. Gertner and Bonnie L. Webber, exemplifies how AI researchers are supporting physicians in their evidence-based

reasoning. This system represents the redesign of a goal-directed planning system for patient management in a trauma unit. Early tests with the original system indicated that working with a dictated plan frustrated the physicians: it contained either too much information (things the physicians already knew) or too little information (for example, no indication of how to improve the plan). In response to this feedback, Gertner and Webber adopted a critiquing approach that addresses the physician's needs by detecting problems with the physician's intended management plan as they arise, and presenting recommendations in the context of the physician's intended actions.

The two application articles present radically different approaches to the production of the primary-care physician's report. Both share two basic goals:

- Make the reporting activity more effective—easy to use, less time-consuming, and more thorough.
- Produce a structured report that the physician can understand and that other computerized systems can use—for data aggregation or longitudinal studies of a patient, for example.

In "Voice-Enabled, Structured Medical Reporting" (pp. 70-73), David F. Rosenthal and Rachael Sokolowski show how they built a tool for physicians. Their tool meets the first goal by providing automatic speech recognition, so the physician doesn't need to type. It achieves the second goal by using SGML and Corba as the underlying standards for data and document structuring.

Alternatively, in "Support for Primary Care Interviews" (pp. 40-48), James R. Warren argues that much of the initial history-gathering activity does not have to be done by a medical expert such as the physician or nurse. Rather, it can be done by a clerk at the front desk who is teamed with an intelligent computerized assistant. In this case, the AI system provides the questions to be answered, and the clerk reads them to the patient and types in the patient's responses. A natural-language-processing component handles some of the more complex responses, and physical systems (such as the respiratory or digestive system) are reviewed according to the patient's answers. The physician is then given this report preparatory to seeing the patient. Preliminary results show that this system compares favorably to physicians' typical reporting methods and has received a good degree of patient acceptance.

2. E.H. Shortliffe, "The Adolescence of AI in Medicine: Will the Field Come of Age in the 90's?" *Artificial Intelligence in Medicine*, Vol. 5, 1993, pp. 93-106.
3. J. Flower, "The Future of Health care," in *Encyclopedia of the Future*, Macmillan & Co., Indianapolis, Ind., 1995; http://www.well.com/user/bbear/health_care_future.html.
4. D.E. Wilson, "Overview and Introduction," in *2020 Vision: Health in the 21st Century (Proc. Inst. of Medicine 25th Anniversary Symp.)*, Nat'l Academy Press, Washington, D.C., 1996, pp. 1-3.
5. J. Hughes, "Doctors Are Dinosaurs in a High Tech World," *New York Newsday*, July 1, 1995; <http://www.changesurfer.com/Hlth/dinosaurs.html>.
6. D.E. Detmer, "Information and Communications," in *2020 Vision: Health in the 21st Century*, Nat'l Academy Press, 1996, pp. 24-36.
7. *Highway to Health: Transforming U.S. Health Care in the Information Age*, Council on Competitiveness, Washington, D.C., 1996; http://nii.nist.gov/pubs/coc_hghwy_to_hlth/tile_page.html.
8. *Cybermedicine: The Computer as a Patient's Assistant* (brochure for a course), Harvard Medical School, Dept. of Continuing Education, Cambridge, Mass., 1997; <http://www.feltco.com/hmscme/data/71011.htm>.
9. J. Flower, *HOTS: Health-Oriented Telecommunications*, 1994, <http://www.well.com/user/bbear/digihealth.html>. (Also published as "The Other Revolution in Health care," *Wired*, Vol. 2.01, Jan. 1994.)
10. D.L. Sackett et al., *Evidence-Based Medicine: What It Is and What It Isn't*, <http://cebml.jr2.ox.ac.uk/ebmisint.html>. (Based on an editorial from *British Medical J.*, Vol. 312, Jan. 13, 1996, pp. 71-72.)
11. Evidence Based Informatics Project, *Evidence-Based Medicine*, Evidence Based Medicine Working Group, Hamilton, Canada, <http://hiru.mcmaster.ca/ebm/overview.htm>, 1997.
12. *Opening* (of the 1996 Howard Paul Medical Symp.), <http://innovate.si.edu/symposia/msbody.htm>, 1996.