Report 83-13
Stanford -- KSL

The Science of Biomedical Computing.
Edward H. Shortliffe,
Mar 1983

Reprinted with permission from the Proceedings of the IMIA International Working Conference on Information Science and Medical Education, J. Anderson, F. Gremy, A. Levy, J.C. Pages, Eds., by North Holland, 1983.

HPP-83-13

The Science of Biomedical Computing

Edward H. Shortliffe

Heuristic Programming Project

Departments of Medicine and Computer Science

Stanford University School of Medicine

Stanford, California 94305

To appear in Proceedings of IMIA International Working Conference on

Information Science and Medical Education

Chamonix, France, March 21-25, 1983.

Edited by J. Anderson, F. Gremy, A. Levy, and J.C. Pages

The author is a Henry J. Kaiser Family Foundation Faculty Scholar in General Internal Medicine and recipient of research career development award LM00048 from the National Library of Medicine. Additional support was provided by the NLM under program project grant LM03395.

1. Introduction

This is a remarkably exciting time to be involved professionally in the field of medical informatics. The underlying scientific principles are beginning to be identified and defined, educators are increasingly acknowledging the importance of the field for physicians of the present and future, and the technology itself is growing at rates that make the future of the field both unbounded and impossible to predict. One has the sense that what was once a field for pioneers is now reaching the stage of established settlements, with a history, traditions, and a feel of permanence.

It is therefore appropriate that, at the beginning of deliberations designed to achieve significant educational goals for the field, we might start by considering the discipline itself and the characteristics that have tended to separate it from other traditional academic and research medical specialties. I would like to begin by assuming that certain basic points are well accepted and need not be defended here: first that medical informatics holds both realized and potential importance for the science and practice of medicine, and second, that there is a need for all medical practitioners to be familiar both with information handling technology and with the underlying principles that make the field relevant, regardless of whether computers are involved. It is accordingly my assumption that conference attendees agree that it is imperative to introduce more effectively the field of informatics into medical education, to help encourage the development and acceptance of medical informatics as a medical science, and to use evolving hardware technology optimally to further these causes.

I should preface my remarks by noting that my observations are largely colored by my familiarity with the situation in North America. I am extremely impressed by the success that many of our European colleagues have had in introducing the field into medical school curricula and in gaining recognition of informatics as an integral part of the quality education of physicians and other health care professionals. The very name "medical informatics" continues to be poorly accepted in the United States, even within the field of medical computer science from which much of the information science work has sprung. The term remains colored with interpretations tied more to library science thán to the more general issues of information management and decision making.

With these introductory remarks and caveats behind me, I would like to focus on the need to define a scientific tradition for the field, one that is akin to other medical disciplines such as biochemistry, genetics, or biophysics. It is my conviction that the failure to define a scientific basis for the emerging discipline has limited its impact on scientists in other medical fields and has hindered its timely introduction into medical education. There are early signs that medical educators are beginning to recognize the field as encompassing a set of fundamental skills needed by all physicians; for example, an ongoing study of medical education by the American Association of Medical Colleges has identified medical information science as a key area for curriculum development. However, there is an inevitable barrier to effectively meeting this challenge -- we do not have the people to teach the courses, to carry out the research, or to serve on an equal footing with other members of medical faculties. Let us consider some of the reasons for a lack of adequately trained people in the field. Then I will give some examples of work in medical informatics that I believe has provided significant new knowledge both for medicine and for computer science. I will close with some specific suggestions for meeting the challenges that lie ahead.

2. The Need for Medical Information Scientists

There is little doubt that we need to attract more bright young minds to the field of academic medical informatics. We have not been optimally successful in these efforts, partly because there is a frequent perception that it is extremely difficult to prepare for or maintain a research career in medical information science. In August of 1982, Dr. Scott Blois and I chaired a small working group that met to discuss mechanisms for encouraging talented young people to consider career paths in the medical informatics field. This group met at the University of Pittsburgh and was organized as one of several parallel sessions at the Annual Workshop on Artificial Intelligence in Medicine. Although a diversity of opinions was represented, remarkable uniformity of expression was achieved in the final report to which I would like to refer here.

We began by asking "What are the principal barriers to producing high quality medical computing researchers and maintaining their commitment to the field?" After considerable discussion, the answers to this question were condensed into seven major areas, presented here in what was judged to be their decreasing order of importance. Once again I should acknowledge that the responses to this question and their relative ordering might be very different in countries other than the United States and Canada.

- Problem #1: Research in medical computer science is constrained and rendered unstable by insufficient and insecure funding; this also hampers the field's ability to compete for trainees and career researchers.

- Problem #2: The efforts of medical computing researchers are constrained by the lack of a recognized niche within their department, their school, the medical and computer science communities, and funding agencies.

- Problem #3: Establishing a credible research career in medical computer science is difficult because start-up costs are high, experiments are lengthy, validation and transfer of technology are complex, and original publications are infrequent and difficult to place in prestigious journals.

- Problem #4: Training is unusually lengthy and requires expertise in two disparate fields, frequently with barriers to cross-departmental cooperation.

- Problem #5: Training is constrained by insufficient funds to support pre- and post-doctoral students, faculty, and computational tools.

- Problem #6: Training is constrained by insufficient curricular opportunities (at both medical and graduate student levels) as well as limited access to advanced computing facilities.

- Problem #7: Medical computer science is a young, interdisciplinary field; thus there are relatively few established investigators to serve as role models and to help define academic career pathways.

The solutions proposed for these problems were numerous and cannot be fully discussed here. However, certain key points pervaded the analysis:

- Those in the field should work to encourage increased recognition of the parallels between medical informatics research and traditional medical research fields;

- Those in the field should work for the establishment of additional medical information science training programs, ideally with customized curricula and graduate degrees when possible;

- Those in the field should work to develop interdisciplinary ties between medical schools and computer science departments (e.g., through interdisciplinary training programs and joint appointments for key faculty).

As I intimated earlier, the first of these items is central to all topics concerning an increasing educational role for medical informatics, and I would therefore like to concentrate on it for the remainder of my talk.

3. The Spectrum of Medical Informatics Activities

Medical computing is, by its nature, *experimental computer science*. An experimental science is characterized by posing questions, designing experiments, performing rigorous analyses, and promoting the feedback of this experience in the design of new experiments. One purpose of such efforts is simply the search for new knowledge. There is a continuum between this fundamental search for new knowledge (so-called "basic research") and the formal application of that knowledge to a specific practical problem ("applications research" or, after a technology is well-developed, "engineering").

Experimental Computer Science

<-->

Basic Research   Applications Research   Engineering

There has been a common assumption that medical computing, because the computer is *applied* to medicine, is *applications research* and involves the use of well-established techniques from the field of computer science. Of course, much of medical computing does fall into this category. It should be noted, however, that other fields of medical research do not use this criterion for distinguishing basic from applied investigation. Throughout medical research there are fields where the work is clearly "basic", even though the work is "applied" to medicine (e.g., investigation into fundamental questions about subjects such as physiology, biochemistry, or immunology). Occasionally a breakthrough in a pure research environment leads to a leap in knowledge that allows steps to be taken in the applied research world. We have seen this happen recently in genetic engineering, where the basic research laboratories, working on limited, well-specified problems, produced some inroads that suddenly spawned a whole new applied research community.

4. The Science Of Medical Computing

It can be argued that medical informatics is still in a "prescientific" era. We build artifacts, our papers describe them, but an underlying theory has thus far failed to emerge. In the field of pure computer science, courses in the mathematical theory of computation exist, but even so, much of that field is characterized by individual experiments and reports of their success or failure. The field of medical informatics has underlying fundamental principles, but we tend to identify and articulate them poorly, and those basic science issues that exist are not generally recognized or accepted. Rarely are scientific activities in the field viewed on an equal level with traditional "pure science" medical research.

Recognition of the youth of medical computing helps put the discipline in perspective when it is compared with other medical research domains. The field is less than a quarter century old, accounting in large part for its prescientific character, and computer science itself has only recently joined the faculties at most of our universities. It is folly to attempt to predict what the field will produce in the future, particularly at a time when hardware advances are revolutionizing many of the basic assumptions that have constrained software design during the last two decades. Perhaps an analogy would clarify this point. Imagine if early radiologists in 1920, 25 years after Roentgen's discovery of x-rays, had tried to predict the technological advances in their field over the next 60 years -- the clarity of images that would one day be achieved, computerized tomography, or the rise of ultrasound and nuclear medicine imaging techniques. Similarly, basic computing research, some of it specifically arising from work in medical domains, will change many of our approaches to clinical systems development in unpredictable ways.

Both observers of the field *and* medical informatics researchers also frequently fail to recognize the link between fundamental clinical issues and medical computing topics. Computing research can lead to important insights regarding issues in medical care, independent of whether the computer is ultimately used. For example, computer-based studies of clinical cognition, by providing insights into the knowledge structures and procedures used by expert practitioners, may facilitate improved training of medical students and result in more thoughtful physicians. Specific problem solving techniques and analytical methods have not traditionally been formally taught, and students have generally learned them through practice and observation of their mentors.

Another reason that the scientific components in medical computing are frequently overlooked is the computer's reputation as a research tool used in the laboratory, rather than as a subject itself worthy of theoretical study. To some observers, computers seem divorced from basic science and are not recognized as potential sources of new knowledge. For example, I recently participated in the review of candidates for a government-sponsored medical scientist training program at Stanford. Most applicants are interested in laboratory research, but one had also expressed a fascination with medical uses of computers. The faculty member evaluating his application observed that, although the student was interested in computers, he was also experienced in laboratory research techniques. The faculty evaluator was therefore relatively sure that the applicant would do "real science" if he came to our University and not get side-tracked by computing issues. I was chagrined but not surprised by the attitude revealed in this comment.

Another characteristic of medical informatics also deserves comment: the length of time required for an experiment. A good experiment in medical computer science, especially if the desired endpoint is clinical use of the new system, typically takes time on the order of years, whereas the wet-bench researcher often can do one or two good experiments, each warranting publication, in a matter of months. Medical faculty who are interested in academic computing are thus presented with conflicting concerns; they must balance the pressure to publish frequent research reports with their desire to describe their work in the literature only when the results are solid and worthy of presentation and defense.

5. Some Examples

I fear that my discussion has concentrated so much on general issues that its persuasiveness may be limited by a lack of specific examples. Therefore I would like to draw briefly on two pieces of current work that I believe have provided major fundamental insights for the field of medical informatics. It happens that both examples are drawn from the field of medical artificial intelligence (AI), the research area with which I am personally most familiar. I am certain, however, that similar excellent examples could be drawn from the other subfields of medical informatics such as medical databases, signal analysis, or image processing.

The first example is the work of Pople, Myers, and Miller in developing the INTERNIST-1 and CADUCEUS programs. This massive ten-year effort has resulted in a large and impressive program that does an admirable job diagnosing single and multiple disorders from among 600 diseases known to the system [6]. The program provides solutions to many problems not tackled in previous systems: an ability to diagnose multiple coincident diseases, an ability to ignore findings that appear to be "red herrings" of low importance, and a hypothesis-directed reasoning method that simulates in many respects the approach to clinical problem solving observed in expert practitioners [4, 5]. The techniques developed are themselves major contributions and have influenced subsequent work in computer-based medical decision making. I have been particularly impressed, however, by the insights into INTERNIST-1's inadequacies that have been learned through the work and thoughtfully discussed by Prof. Pople in a recent book on medical AI edited by Prof. Szolovits of MIT [7]. It is important to emphasize that the lessons could not have been learned if INTERNIST-1 had not been built. Thus the work is in the best traditions of experimental science: early work, often arduous, leads to results and insights which define the next iteration in the process. Among the key points cited in Pople's chapter is the need for a decision making system that does the following: (a) deals with a "tangled network" of hierarchic knowledge reflecting the nosology of both organ system involvement and etiology, (b) encodes and uses the basic causal knowledge that allows reasoning from first principles and the dynamic recognition of syndromes and interrelated disease states, and (c) uses dynamic task synthesis operations which aid in the planning of the diagnostic approach. These insights, and their subsequent impetus for the development of the CADUCEUS program (the successor to INTERNIST-1), are a tribute to the early research and demonstrate why it is a fallacy to call a medical informatics experiment a failure if it is not implemented for clinical use. This criterion is not used in assessing other fundamental medical research, and its frequent use in judging medical computing work is a reflection of the failure to appreciate the basic science issues in our field and of the unrealistic expectations that frequently exist.

The second example I would like to mention is drawn from the field of medical computer-aided instruction and is therefore particularly pertinent for discussion at this meeting. The program is the GUIDON work of Clancey [1], a system designed to teach the knowledge in the MYCIN program. Our earlier MYCIN work at Stanford had focussed on producing a high performance consultation system that could advise physicians in the selection of antibiotic therapy [8]. Most of the knowledge in this domain was encoded in production rules and Clancey suggested that a useful project would be to devise an interactive program that could teach these rules to medical students. The resulting program, GUIDON, was an excellent and provocative piece of work, but as was true for INTERNIST-1, the insights regarding its inadequacies were in some ways more important than the actual level of performance it achieved. The lessons learned were largely unexpected when the effort began. For example, Clancey learned that rules that allow a program like MYCIN to reach good decisions are often opaque when used in a teaching environment. The underlying "support knowledge" which explains *why* the rule is true was missing from MYCIN and yet crucial for good instruction of students. Similarly it became clear that simple memorization of rules was an inadequate way to learn about a complex clinical domain. Equally important was instruction in the strategies for effective problem solving, techniques for invoking or discarding potentially useful background knowledge, and effective guidance in how to evoke hypotheses and then constrain the differential diagnosis [3]. After extensive studies of excellent teachers interacting with medical students, Clancey identified the need to encode and use strategic knowledge as well as support knowledge in a system like MYCIN and he developed a revised version of the system known as NEOMYCIN [2]. This has in turn formed the basis for his current work on an improved teaching program known as GUIDON-2.

Both these examples reflect the experimental nature of fundamental medical informatics research and illustrate the iterative process that tends to drive such investigations. Particularly important has been the investigators' efforts to articulate the insights gained and to motivate their subsequent work accordingly. I will return to this point later in the discussion.

6. Some Recommendations for the Future

As we look ahead to the next decade of medical informatics research, it is useful to look for ways to guide the field from a prescientific to a scientific era. The environment for this progression should be excellent: departments of medical information sciences are growing in number and reputation, especially in Europe; increasing numbers of researchers, sometimes with formal training in both medicine and computer science, are participating in investigative work; the formal degree programs are resulting in innovative high quality research by PhD candidates; the federal government in the United States is supporting (albeit somewhat tenuously at present) formal education in the field through training grants; and there is a growing interest in computing among practitioners, as witnessed by the recent flurry of continuing education courses regarding programming and office computing.

The following suggestions may provide some useful guidelines for encouraging understanding and acceptance of the scientific elements of the field and for heightening the field's role in medical education:

- Medical information scientists should more specifically identify the basic science components of ongoing research projects. This means a careful analysis of the task domain when research is proposed, and an assessment of which components of the task, if any, cannot be easily handled by state-of-the-art techniques. Applications work requires scrupulous attention to task selection and interface specification if finished, accepted products are desired in the short term. Research proposals must make it clear which parts of a project will require innovative techniques and are therefore potential stumbling blocks in the smooth progression of a project.

- In addition to the introduction of medical information science training into medical schools, there is a need for well-designed continuing education courses for health personnel. All such courses must provide realistic appraisals of the field, including frank characterization of the current limitations, specification of the research problems currently under attack, and cautious optimism about the future of medical computing and its considerable promise for assisting with the organization and delivery of modern health care. Wasserman has provided a pertinent list of important problems for medical computer science [9]; many of these can be conveyed to any health professional and need not be saved for degree candidates in formal medical information science training programs.

- Formal papers and research presentations should discuss the basic components of the work being reported. It is much more likely that colleagues will be able to build on an investigator's work if the research themes and new insights are emphasized, rather than high-level descriptions of a specific new computing project.

I would like to close by focusing on this third point in somewhat more detail. Although any generalization invites counter-examples, and I know there are many that could be cited in this case, I would like to argue that the quality of the medical informatics literature is one of the most serious impediments to the field's emergence as a scientific discipline. This point was recently emphasized for me when a medical student I know undertook an extensive review of the literature on medical database systems. In his report on this study, he made the following observation:

> Medical system implementors don't read (presumably because they "know" what they want). An enormous duplication of effort, both good and bad takes place. Bibliographies are typically incomplete, and fundamental issues are almost always overlooked. When reading is done, it is done incestuously.

Although I have heard similar sentiments expressed previously, and tend to agree with them, it was particularly impressive to me that a bright student of medicine, one who is familiar with articles in other areas of medical science, should perceive so quickly and insightfully this underlying problem in the literature of medical informatics.

What is the answer, then? How do we improve the quality of the reports that appear in our journals and conference proceedings? I would suggest two approaches: the development of improved guidelines (in Instructions to Authors) for the format of medical informatics articles, and the implementation of more rigorous review processes before articles are accepted for publication.

Regarding the first of these ideas, many of us have experienced the problem of trying to write a paper about medical informatics research when a journal recommends the traditional "Introduction, Methods and Materials, Results, Discussion, and Conclusions" organization. Except for formal reports describing a clinical trial of a well-developed system (e.g., one in which the statistical design of the study constitutes a suitable topic for a "Methods" section), papers in medical informatics generally do not fit naturally into the usual outline. The following alternative organization for a medical informatics report no doubt requires refinement, and it is more appropriate for a project report than for the presentation of a new technique, but it does address many of the issues which account for our medical student's poor opinion of the current literature:

Introduction: a discussion of the clinical issues underlying a project's design, the need for the system, and any pilot studies done to demonstrate the demand for such a program.

Background: this should be divided into two subtopics: a discussion of prior work by the authors that led to the current design decisions, and an analysis of related work in the literature. A mere listing of related work would be inadequate here; an attempt to acknowledge and integrate lessons from the work of others is crucial.

Design Considerations: a discussion of the principal design and implementation goals against which the success of the system should be assessed.

System Description: a technical exposition of the overall architecture, the pertinent data structures, control mechanisms, etc. A brief description of the hardware used may be useful here but this is seldom worthy of extensive discussion in a paper from which underlying principles are to be derived.

Status Report: a description of the current status of the implementation, informal indicators of the program's strengths or weaknesses, examples of its current level of performance.

Lessons Learned: an analysis of the key insights gained from the work to date, focussing on the statement of general principles that can contribute to the knowledge of the field. The statement of such lessons must be well defended with examples. The emphasis should be on the statement of principles in such a form that they will be of use to other investigators in the field.

Future Plans: a discussion of how the research to date, and the lessons learned, have led to a specification of future research goals and possible revisions in the program's design or implementation.

One advantage of such an organization of reports is that it provides a criterion for determining when a research group is ready to publish a paper. If there is nothing to say under "Lessons Learned", one might argue that it would be better to wait.

I doubt, however, that such requirements would have a drastic negative impact on the number of papers submitted or published. It is my impression, rather, that we simply tend not to communicate well the lessons learned in our research. General descriptions of systems are frequent, but the definition of underlying principles or insights are simply not yet a regular part of the routine. If reviewers began to demand such discussions, the awareness of the need for them would be heightened and I suspect we would see a gradual change in the quality of the papers, and perhaps in the research itself.

I do not doubt that there exist alternate views on an optimal organization for medical informatics papers, and I do not mean to suggest dogmatically that one outline should suffice for all reports in the field. However, the central issue - an attempt to extract the fundamental principles of our discipline - will in my opinion be an element of any effective solution to the problem. Such efforts will accordingly contribute to the goal that brings us together this week: the introduction and acceptance of medical informatics as an established discipline with a central role in the education of health professionals.

References

1. Clancey, W.J. *Transfer of rule-based expertise through a tutorial dialogue*. Ph.D. Th., Stanford University, September 1979.

2. Clancey, W.J. and Letsinger, R. NEOMYCIN: Reconfiguring a rule-based expert system for application to teaching. , 7th International Joint Conference on Artificial Intelligence, Vancouver, British Columbia, August, 1981.

3. Clancey, W.J. "The epistemology of a rule-based expert system: a framework for explanation." *Artificial Intelligence* (1983). in press

4. Elstein, A.S., Shulman, L.S. and Sprafka, S.A.. *Medical Problem Solving: An Analysis of Clinical Reasoning.* Harvard Univ. Press, Chambridge, MA, 1978.

5. Kassirer, J.P. and Gorry, G.A. "Clinical problem solving: A behavioral analysis." *Ann. Intern. Med 89* (1978), 245-255.

6. Miller, R.A., Pople, H.E., Myers, J.D. "INTERNIST-1, An experimental computer-based diagnostic consultant for general internal medicine." *New England Journal of Medicine 307* (1982), 468-476.

7. Pople,H.E., Jr. . In *Artificial Intelligence in Medicine*, Szolovits, P., Ed.,Westview Press, Inc., Boulder, Colorado, 1982, ch. Heuristic methods for imposing structure on ill-structured problems: The structuring of medical diagnostics, pp. 119-190.

8. Shortliffe, E.H.. *Computer-Based Medical Consultations: MYCIN.* Elsevier/North Holland, New York, 1976.

9. Wasserman, A.I. Educational programs in medical computing: their basis and implementation. , 3rd Symposium on Computer Applications in Medical Care, Silver Spring, Md., October, 1979, pp. 217-222.