Report 84-20
Stanford -- KSL

Artificial Intelligence and the Future of Medical Computing.

Edward H. Shortliffe,

Feb 1984

Heuristic Programming Project
Report No. HPP 84-20

February 1984

Artificial Intelligence and the Future of Medical Computing

Edward H. Shortliffe
Stanford University, Medical School

Published in the Proceedings of a Symposium on Computers in Medicine, Annual Meeting, California Medical Association, Anaheim, CA., February 1984.

Artificial Intelligence and the Future of Medical Computing

Edward H. Shortliffe, M.D., Ph.D.
Division of General Internal Medicine, Department of Medicine
Stanford University School of Medicine
Stanford, California 94305

Although computing technology is playing an increasingly important role in medicine, systems designed to advise physicians on diagnosis or therapy selection have remained largely experimental to date. Despite diverse research efforts, and a literature on computer-aided diagnosis that has numbered over 1500 references in the last 20 years, clinical consultation programs have failed to achieve wide acceptance.

The reasons for attempting to develop such systems are self-evident. Growth in medical knowledge has far surpassed the ability of the single practitioner to master it all, and the computer's superior information processing capacity thereby offers a natural appeal. Furthermore, the reasoning processes of medical experts are poorly understood; attempts to model expert decision making necessarily require a degree of introspection and a structured experimentation that may in turn improve the quality of the physician's own clinical decisions, making them more reproducible and defensible. New insights that result may also allow us more adequately to teach medical students and house staff the techniques for reaching good decisions, rather than merely to offer a collection of facts which they must independently learn to use coherently.

In recent years observers have begun to analyze the reasons for poor acceptance of the systems that have sprung from such research [2]. We have argued [6, 8, 9] that the problems have tended to lie not only with the decision-making performance of such programs but also with system design features that have failed to appreciate the physician's viewpoint or have made the interactive process unappealing. To correct these deficiencies, future systems must be fast, easy to use, and congenial. They must address important clinical problems with which physicians recognize they need assistance. But perhaps most important, in order to stress the primary physician's role as ultimate decision maker, they must be able to explain what they are doing, not through quotations of statistical theory but in terms of a line of reasoning that is familiar and similar to the kind of justification a clinician might expect from a human consultant.

Explanation capabilities help the physician who is using the program decide whether to follow its advice; they thereby emphasize the computer's function as a helpful tool that complements rather than replaces the primary physician's own decision-making powers.

Because of considerations such as these, the last decade has witnessed the development of new approaches to computer-based medical decision making. Of particular significance is research directed at the encoding and use of experts' judgmental knowledge -- the kind of practical experience which underlies the daily practice of medicine and is far-removed from the mathematical approaches of formal statistical analysis. One relevant computer science subfield, termed "artificial intelligence" because of its emphasis on symbolic reasoning capabilities rather than numeric computations, has formed the theoretic foundation for several of these experimental consultation programs [10]. Some of the early work emphasized excellent decision making performance combined with an ability to explain the basis for the program's advice. The MYCIN program developed at Stanford was one such system. It was an experimental program that assisted with the selection of antimicrobial therapy for patients with infections [5]. Knowledge of bacteremia and meningitis was acquired from infectious disease experts and encoded in decision "rules". This knowledge was in turn used by a program which considered a specific case, interacting with the person requesting advice and generating a therapeutic recommendation. By responding to specific questions asked by the physician and translating computer code into understandable English, MYCIN was also able to explain the basis for its decisions so that the user could independently decide whether to follow its advice. In addition, MYCIN was demonstrated to function similarly to infectious disease experts when selecting therapy for isolated bacteremias [11] or meningitis [12]. A variety of logistical issues, including the size and expense of the computers necessary to run programs of this complexity, prevented the program's implementation for routine clinical use. Recent hardware developments, with machines that are continually more powerful becoming available in small desktop instruments for greatly reduced prices, will make the introduction of large and complex advice systems such as MYCIN more realistic in the coming years.

A second medical artificial intelligence system of considerable promise is the INTERNIST/CADUCEUS program developed at the University of Pittsburgh [3]. This large program has focussed on the diagnostic task, encoding knowledge about the manifestations of more than 600 diseases in internal medicine and neurology. The program's ability to reach accurate conclusions about complex cases such as those appearing as CPC's in the clinical literature has demonstrated the potential in this kind of approach. However, as the developers have cogently explained in a recent article [4], major theoretical and logistical problems remain to be solved before a system suitable for routine clinical use will be available.

For more narrow problems, however, we are now beginning to see systems coming into routine use in University hospitals. At Stanford University Medical Center, for example, the ONCOCIN program has been used routinely for over two years in the management of cancer patients enrolled in chemotherapy protocols and undergoing treatment for any of a variety of lymphomas [1]. This system uses artificial intelligence techniques to encode the knowledge of the expert oncologists who designed the protocols and who advise physicians regarding their proper application. However, much of the system design effort has focussed on the development of optimal techniques for integrating this kind of an advice system into the clinical environment in such a way that physicians find it helpful, easy to learn and use, non-threatening, and respectful of their busy schedules. Ongoing efforts to transfer ONCOCIN to run on single-user small computers will soon make it possible to move this consultation/data-management system beyond the university environment to private offices and small clinics.

These kinds of promising research into medical symbolic reasoning represent more than the application of well-established computing techniques. The approaches are young, experimental, and they demand refinement. Extensive basic medical computing investigation will be required, in artificial intelligence and in related computer science subfields, before useful, congenial, high-performance consultation systems will be both available <u>and</u> accepted by physicians. Effective research in this area will also require a new generation of scientists with extensive training both in medicine and in computer science. Innovative programs to train such individuals are beginning to appear at a handful of U.S. medical and engineering schools.

Successful systems are unlikely until we know better how to manage such problems as:

- understanding the psychology of medical reasoning as practiced by specialists,
- automating the interpretation of written and spoken natural language,
- developing new uses of the graphics capabilities of modern computers, both to simplify the interaction with the physician and to experiment with new modes of explanation and teaching,
- devising simpler approaches to the acquisition and representation of knowledge obtained from collaborating experts,
- encoding and using the temporal relationships that are central to many disease processes,

- representing and measuring inexact reasoning.

The potential of such research is great, but progress may depend in part on increased recognition that computer science represents an important form of basic medical investigation. Improved understanding of these issues will depend on improved educational opportunities in medical information sciences, both for medical students and for physicians already on faculties or in private practice settings.

This discussion is an updated and expanded version of an editorial that originally appeared in the Archives of Internal Medicine [7].

References

1. Bischoff, M.B., Shortliffe, E.H., Scott, A.C., Carlson, R.W., and Jacobs, C.D. Integration of a computer-based consultant into the clinical setting. Proceedings of the 7th Symposium on Computer Applications in Medical Care, IEEE, Baltimore, MD, 1983.

2. Friedman, R.B. and Gustafson, D.H. "Computers in clinical medicine: a critical review." *Comput. Biomed. Res. 8* (1977), 199-204.

3. Miller, R.A., Pople, H.E., and Myers, J.D. "INTERNIST-1: An experimental computer-based diagnostic consultant for general internal medicine." *New Eng. J. Med. 307* (1982), 468-476.

4. Pople, H.E. Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnostics. In *Artificial Intelligence in Medicine*, Szolovits, P., Ed.,Westview Press, Boulder, CO, 1982, ch. 5, pp. 119-190.

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

6. Shortliffe, E.H., Buchanan, B.G., and Feigenbaum, E.A. "Knowledge engineering for clinical decision making: a review of computer-based clinical decision aids." *Proceedings of the IEEE 67* (1979), 1207-1224.

7. Shortliffe, E.H. "The computer as clinical consultant." *Arch. Int. Med. 140* (1980), 313-314.

8. Shortliffe, E.H. Computer-based clinical decision aids: some practical considerations. Proceedings of the AMIA Congress 82, American Medical Informatics Association, San Francisco, CA, 1982, pp. 295-298.

9. Shortliffe, E.H. "The computer and medical decision making: good advice is not enough." *IEEE Engineering in Medicine and Biology Magazine 3* (1982), 24-27.

10. Szolovits, P. (editor). *Artificial Intelligence in Medicine*. Westview Press, Boulder, CO, 1982.

11. Yu, V.L., Buchanan, B.G., Shortliffe, E.H., et al. "An evaluation of the performance of a computer-based consultant." *Comput. Prog. Biomed. 9* (1979), 95-102.

12. Yu, V.L., Fagan, L.M., Wraith, S.M., et al. "Antimicrobial selection by a computerized consultant: a blinded evaluation by infectious disease experts." *J. Amer. Med. Ass. 242* (1979), 1279-1282.