Report 81-07
Stanford -- KSL

Scientific DataLink

ONCOCIN: An Expert System for Oncology Protocol Management.
Edward H. Shortliffe, A. Carlisle Scott, Miriam B. Bischoff, et al., Jul 1981

card 1 of 1

Memo HPP-81-7
Stanford Heuristic Programming Project

ONCOCIN: AN EXPERT SYSTEM FOR ONCOLOGY PROTOCOL MANAGEMENT

Edward H. Shortliffe

A. Carlisle Scott

Miriam B. Bischoff

A. Bruce Campbell

William Van Melle

Charlotte D. Jacobs

Heuristic Programming Project

Departments of Medicine and Computer Science

Stanford University, Stanford, California 94305

July 1981

To appear in:

Proceedings of the 7th IJCAI

24-28 August 1981

Vancouver, British Columbia

This work was supported by the National Library of Medicine under Program Project Grant LM-03395 and research career development award LM-00048. The program was developed on the SUMEX Computer Project at Stanford University, a shared national resource supported by the Biotechnology Resources Program of the NIH under grant RR-00785.

Abstract

We describe an oncology protocol management system, named ONCOCIN, that is designed to assist physicians in the treatment of cancer patients. The system is a set of related programs, one of which is a rule-based reasoner that encompasses the necessary knowledge of cancer chemotherapy. Representation and control techniques are discussed, and ONCOCIN is contrasted with systems that could be built using EMYCIN. Of particular interest is the need to provide ONCOCIN with an interface that will make the system acceptable to oncologists.

Acknowledgments

The ONCOCIN System described in this report has relied on the efforts of several physicians and computer scientists. We would like to express our gratitude to the following present and past project members: Paul Chang, Jai Cho, Chuck Clanton, Larry Fagan, Phil Gerring, Patricia Pickett, Brandy Sikic, Motoi Suwa, Randy Teach, and Craig Tovey.

1 Introduction

This report describes an oncology protocol management system, named ONCOCIN after its domain of expertise (cancer therapy) and its historical debt to MYCIN [5]. The program consists of a set of interrelated subsystems, the principal ones being:

(1) the Reasoner, a rule-based expert consultant that is the core of the system and the major subject of this report; and

(2) the Interviewer, an interface program that controls a high-speed terminal and the interaction with the physicians using the system.

Work on ONCOCIN began in mid-1979, and the system was installed for prelimianry use in May 1981. This paper describes only the Reasoner, although other system components are mentioned when their interface with the Reasoner is pertinent. We also contrast ONCOCIN with EMYCIN [8], explaining why the EMYCIN formalism was inadequate for our purposes although it did strongly influence the system's rule-based design.

2 Overview of the Problem Domain

ONCOCIN is designed to assist clinical oncologists in the treatment of cancer patients. Because the optimal therapy for most cancers is not yet known, clinical oncology research is commonly based on formal experiments that compare the therapeutic benefits and side effects (toxicity) of proposed alternative disease treatments. "Cancer" is a general term for many diseases

having different prognoses and natural histories. A treatment that is effective against one tumor may be ineffective against another. Thus a typical cancer research center may conduct many simultaneous experiments, each concerned with a different kind of cancer and its optimal therapy (i.e., the treatment plan with the best chance of cure, remission, or reduction in tumor size, and the least chance of serious side effects).

Each of these experiments is termed a "protocol". Patients with tumors of the type being studied, and who are accepted for protocol treatment, are randomly assigned to receive one of two or more possible treatments. The experiment requires close monitoring of each patient's clinical response and treatment toxicity. These data are tallied for all patients treated under the alternative regimens, and in this way the "state-of-the-art" is updated over time.

Each protocol is described in a detailed document, often 40 to 60 pages in length, which specifies the alternative therapies being compared and the data that need to be collected. No single physician is likely to remember the details in even one of these protocol documents, not to mention the 30 to 60 protocols that may be used in a major cancer center. Although an effort is made to have the documents available in the oncology clinics when patients are being treated for their tumors, it is often the case that a busy clinic schedule, coupled with a complex protocol description, leads a physician to rely on his memory when deciding drug doses or what laboratory tests to order. Furthermore, solutions for all possible treatment problems cannot be spelled out in protocols. Physicians use their own judgment in

treating these patients, resulting in some variability in treatment from patient to patient. Thus patients being treated on a protocol do not always receive therapy in exactly the manner that the experimental design suggests, and the data needed for formal analysis of treatment results are not always completely and accurately collected.

The problems we have described reach far beyond the oncology clinic at Stanford Medical Center. There are now several institutions designing protocol management systems to make the details of treatment protocols readily available to oncologists and to insure that complete and accurate data are collected[1]. ONCOCIN is superficially similar to some of the developing systems, but its research objectives are unique in ways we describe in the following section. One overriding point requires emphasis: in order to achieve its goals, ONCOCIN must be used directly by busy clinicians; the implications of this constraint have pervaded all aspects of the system design.

3 Design Considerations

The design of the ONCOCIN system has been influenced by the two parallel thrusts of the project's research aims. The first is to perform

[1] A memo from the MIT Laboratory for Computer Science [7] describes a recent collaboration between MIT and oncologists who have been building a protocol management system at Boston University [3]. They are planning to develop a program for designing new chemotherapy protocols. To our knowledge, this is the only other project that proposes to use AI techniques in a clinical oncology system. However, the stated goals of that effort are different from those of ONCOCIN.

research into the basic science issues of applied artificial intelligence. Efforts are being made to improve the tools currently available, and to develop new tools, for building knowledge-based expert systems for medical consultation. Other areas of AI research include the extension of methods of interaction between rule-based consultation systems and a large database of time-oriented clinical information, decision making based upon trends over time, and the development of techniques for assessing knowledge base completeness and consistency.

The second thrust is to develop a clinically useful oncology consultation tool. One of the primary goals of the project is to demonstrate that a rule-based consultation system with an explanation capability can be usefully applied in a busy clinical environment. While working to achieve this goal, we hope to establish both an effective relationship with a specific group of physicians, and a scientific foundation, that will together facilitate future research and implementation of computer-based tools for clinical decision making.

4 System Overview

The ONCOCIN system will eventually contain knowledge about most of the protocols in use at the Oncology Clinic at Stanford Medical Center. Although protocol knowledge is largely specified in a written document, many questions arise in translating the information into a computer-based format. Knowledge base development has therefore been dependent on the active

collaboration of Stanford oncologists. We have started by encoding the knowledge contained in the protocols for treatment of Hodgkins Disease and the non-Hodgkins lymphomas[2].

After examining a patient, the physician uses a video display terminal to interact with ONCOCIN's data-acquisition program (the Interviewer). This session normally includes a review of time-oriented data from the patient's previous visits to the clinic, the entering of information regarding the current visit, and the receipt of recommendations generated by the "Reasoner" regarding appropriate therapy and tests. The Reasoner and Interviewer are linked with one another as shown in Fig. 1. In generating its recommendation, the Reasoner uses initial data about the patient's diagnosis, data about previous treatment, results of current laboratory tests, plus the protocol-specific information in its knowledge base. Before terminating an interaction, the physician can examine the explanation provided with each recommendation[3]. The physician may approve or modify ONCOCIN's recommendation; any changes are noted by the system and kept available for future review. ONCOCIN also provides hard-copy backup to complement the on-line interaction and facilitate communication among clinic personnel.

---

[2]We also implemented the complex protocol for treating oat cell carcinoma of the lung. Because the oat cell protocol is the most complex at Stanford, and it took only a month to encode the relevant rules, we are confident that the representation scheme we have devised will be able to manage, with only minor modifications, the other protocols we plan to encode in the future.

[3]We have chosen a representation that will also eventually allow ONCOCIN to offer a justification for any intermediary conclusions that the system made in deriving the advice.

To ensure that busy clinicians will find ONCOCIN fast and easy to use (as well as simple to learn), a special terminal interface has been built. This Interviewer (Fig. 1) includes hardware and software features not often found in current medical systems. A customized keyboard incorporates a keypad for most of the entry and command functions. The Interviewer uses a high-speed video display terminal with multiple windows; simulating the appearance of the flowsheet that currently appears in the clinic chart and which is used to record all the data required for adequate protocol analysis. As the physician enters flowsheet information, relevant patient data are passed to the Reasoner, which is simultaneously considering the case and preparing a recommendation to pass back to the Interviewer for display. The need for a terminal interface that provides rapid response to user type-in, while simultaneously satisfying the computational demands of the reasoning process, has led us to design a complex system architecture with asynchronous processes. The Reasoner and the Interviewer each run in a separate fork under the TENEX or TOPS-20 operating systems, thereby approximating a parallel processing system[4]. Separation of the interface program from the reasoning program allows the user to receive prompt attention when entering flowsheet values; the Interlisp program, running in background, is not permitted to degrade the interaction.

[4]Another program, the Interactor, handles interprocess communication. There is also a process that provides background utility operations such as file backup. The Reasoner is written in Interlisp, the Interviewer in SAIL.