AN ANTIBIOTIC THERAPY SELECTOR WHICH PROVIDES FOR EXPLANATIONS

By William J. Clancey
Heuristic Programming Project
Department of Computer Science
Stanford University
Stanford, California 94305

Experience with the MYCIN [1] therapy selector shows that a cleanly structured algorithm makes possible a si.-nple, but useful explanation capability. The algorithm uses the generate and test method to select a small set of drugs for administration to a patient having an infectious disease. Traces of the application of medical strategies are left behind for later selective retrieval and printing by the explanation system. These strategies are readily comprehensible to the user of the program because they are based on existing clinical rationale for selecting antibiotics (as opposed to using an evaluation polynomial, for example.)

The generate and test algorithm is described as a series of steps through which each recommended drug must proceed successfully:

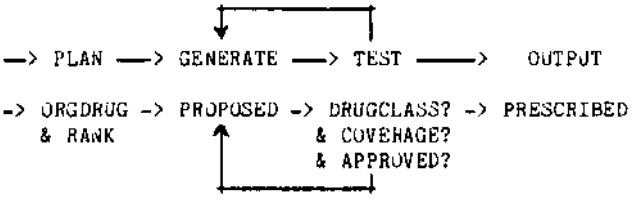

Each step corresponds to a subset operation on surviving candidates from the previous step. The predicate of the subset operation is the collection of medical strategies which characterize the drugs at each step. As shown in the diagram above, these steps are:

ORGDRUG — the drug is one of the drugs which might be prescribed for an organism requiring therapy

RANK — the preference ranking of the drug, as determined by sensitivity information and whether it is a current therapy

PROPOSED — the drug is being considered as part of the regi.men for this patient:

The generator selects subsets of drugs for testing under the control of instructions which specify how many drugs to select from the preference categories, e.g., "propose 2 drugs: one second choice and one third choice." Instructions for regimens containing 1 or 2 drugs are taken from a static list; for regimens containing 3 or more drugs, instructions are generated from a simple pattern. Instructions are ordered to conform to clinical practice for selecting therapy. Thus, two second choice drugs will be proposed before a combination of a second and third chojce drug.

DRUGCLASS? -- does the possible regimen include 2 drugs from the same drug group? (e.g., the penicillins)

COVERAGE? — does the possible regimen cover for all of the organisms that require therapy?

APPROVED? — are there contraindications for this regimen? (e.g., allergic reaction)

PRESCRIBED — the drug is one of the drugs recommended for this patient

This list of steps is maintained separately from the algorithm, and it, along with lists of medical criteria applied at each step, constitute a (static) EVENT STRUCTURE which is used by the explanation system to retrieve traces. The medical criteria are referred to as positive and negative REASONS, because a recital of the chain of positive reasons will serve as an explanation for why a drug was prescribed. Similarly, the occurrence of a negative characterization during the execution of the algorithm will serve as an explanation for why a drug was not prescribed, e.g., the drug was rejected because laboratory tests showed that the organism is resistant to it. The event structure may be thought of as a state-transition diagram of the algorithm, specifying steps (states) and reasons why a drug makes it to the next step (transit ions).

In summary: 1) the algorithm leaves behind a trace indicating which medical criteria were applied to each drug, this constitutes a (dynamic) EVENT HISTORY, and 2) the explanation subprogram retrieves the relevant traces for a particular drug by "reading" the (static) EVENT STRUCTURE of the algorithm which lists the medical criteria which might have been applied at each step. The retrieval process is therefore domain independent, permitting the explanation program to be separable from the knowledge base.

We find that this simple process of retrieving traces works because 1) the number of traces is relatively small (there are fewer than 50 drugs and usually fewer than 100 proposals); 2) the steps of ranking drugs, proposing and testing regimens are readily comprehensible to the user because they attempt to mimic his heuristics for selecting therapy; 3) the explanation subprogram has one basic question: "why was(n't) a particular drug prescribed for a particular organism?", greatly simplifying the organization of traces.

ACKNOWLEDGMENTS: The therapy selector was developed by Victor Yu, M.D., Shari Wraith, B.S. Pharm, and the author. Part of the generator was coded by William van Melle. We benefited greatly from experience with Ted Shortliffe's original algorithm. [Research sponsored in part by Bureau of Health Services Research and Evaluation, Grant HS01544 and ARPA Contract DAHC15-73-C-0435. SUMEX-AIM computer facilities supported by Biotechnology Resources, Grant RR-00785.]

A more detailed working paper with the same title is available from the author.

[1] Shortliffe, E.H. (1976) Computer-Based Medical Consultations: MYCIN, American Elsevier