THE USE OF MODELS IN A RULE-BASED CONSULTATION SYSTEM

Jan Aikins
Heuristic Programming Project
Department of Computer Sci ence
Stanford University
Stanford, California 94305

Research[1] on Mycin-like[2] consultation systems indicates that a rule-based system achieves better performance if guided by models of "typical" consultation situations. One very powerful form of reasoning used by a human consultant when he is presented with a new situation is that of classifying and comparing the new situation with situations he has encountered previously. For example, an auto mechanic who is asked to fix a car which won't start will check the starter, the battery, and the supply of gas because he is guided by his model of a "car won't start" situation, and tnus has expectations about the cause of the problem. The auto mechanic is also able to check the consistency of his information precisely because he has a model of the typical findings for this situation.

Similarly, in our consultation system, we use models of typical consultation situations to guide the invocation of our rules, and to check for inconsistent information. Our models are frame-like[3] structures containing both model-independent pieces of information called elements, such as a name and a certainty measure, and model-specific pieces of information called the components of the model. For example, the components in a model of acute cystitis include the age, sex and therapy for the typical acute cystitis patient.

Inability to fit information into one or more of these models may indicate that there are inconsistencies. For example, finding sickle-cell anemia in a Caucasion patient is not consistent. It may also indicate that the structure of the model is incomplete, for instance, that a component is missing, or that this is an atypical situation that is not represented in any model.

There exists a hierarchy of models, beginning at the top with a model of a general consultation and continuing down to more specific models. For example, models in one general-to-specific hierarchy would be "general consultation", "Mycin consultation", "urinary tract infection", and "acute cystitis". Each model also contains a list of suggested alternates to be tried when that model fails to fit the current situation. For example, an alternate for "acute cystitis" is "chronic cystitis".

At any moment, there is one model which represents the system's current hypothesis about how to classify the given information. Attempting to fill in the components of this model may in turn cause rules to be executed or questions to be asked of the user. Therefore, rules are executed and questions are asked for the purpose of verifying this current hypothesis. In the system without models, rules are executed to establish a general goal, such as, "determine if the patient has disease", which may cause irrelevant rules to be tried. The search for information in a model-based system is thus more directed and faster because irrelevant questions and needless rule invocations are minimized.

Models also allow an improved explanation facility because we can say what the current-hypothesis is, that is, which model is being tried. At present, the explanation facility can only state which rule is being executed, which rule was executed before it, etc., in a process of "unwinding" the rules which were used in making decisions. In the model-based system, rule invocations occur in the context of a model so that a "higher-level", contextual explanation can be given about why these rules are being tried.

The U3e of models gives us further capabilities, such as the ability to generate patient summaries, which have been difficult to achieve previously because this contextual information was not available. In addition, models can be updated or added to the system wi th relative ea3e, allowing us to accomodate a large and changing body of technical knowledge. Thus we feel that the adaptation of the current rule-based formalism to include models results in a substantially improved and very powerful consultation system.

1 This work was supported by the Advanced Research Projects Agency under contract DAHC 15-73-C-0M35. Computer facilities were provided by the SUMEX-AIM facility at Stanford University under National Institutes of Health grant RR-00765. The author is sponsored by the Xerox Corporation under the direction of the Xerox Palo Alto Research Center.

2 See generally, E. Shortliffe, MYCIN: Computer-Based Medical Consultations. New York: American Elsevier, 1976.

3 M. Minsky, "A Framework for Representing Knowledge", in P. Winston (Ed.), The Psychology of Computer Vision, New York: McGraw-Hill, 1975.