

Management of Uncertainty in Medicine

COINS Technical Report 86-12

Paul Cohen, David Day, Jefferson DeLisio,
Michael Greenberg, Rick Kjeldsen, Dan Suthers

Experimental Knowledge Systems Laboratory
Department of Computer and Information Science
University of Massachusetts, Amherst, MA

Paul Berman, M.D.
University of Massachusetts Medical School

Abstract

MUM is a knowledge-based consultation system designed to manage the uncertainty inherent in medical diagnosis. The primary task of the system is to plan which questions, tests, and treatments to order at each point in a consultation given its current uncertain knowledge about the patient's diseases. Managing uncertainty means planning what to do when uncertain; we suggest that this ability must be designed in, not added on, to the architectures of knowledge based systems. MUM is based on one such architecture, implemented as a generalized inference network and planner. The network facilitates local combination of evidence; the planner "reads" the state of the network after each piece of evidence integrated, then decides which evidence to seek based on its several goals.

This research is funded by DARPA/RADC Contract F30602-85-C-0014 and by National Science Foundation Grant IST 8409623.

Published in *International Journal of Approximate Reasoning*, 1986, Vol.1, No. 1, pp. 103-116.

1. Introduction

MUM is a knowledge-based consultation system designed to manage the uncertainty inherent in medical diagnosis (the acronym stands for Management of Uncertainty in Medicine). Managing uncertainty means planning actions to minimize uncertainty or its consequences. Thus it is a control problem – an issue for the component of a knowledge system that decides how to proceed from an uncertain state of a problem. Uncertainty can be managed by many strategies, depending on the kind of problem one is trying to solve. These include asking for evidence, hedging one's bets, deciding arbitrarily and backtracking on failure, diversification or risk-sharing, and worst-case analysis. The facility with which a consultation system such as MUM manages uncertainty is evident in the questions it asks: it should ask all necessary questions, no unnecessary questions, and it should ask its questions in the right order. These conditions, especially the last one, preclude uniform and inflexible control strategies. They prompted the development of the MUM architecture in which control decisions are taken by reasoning about features of evidence and sources of uncertainty.

1.1 The Goals of MUM

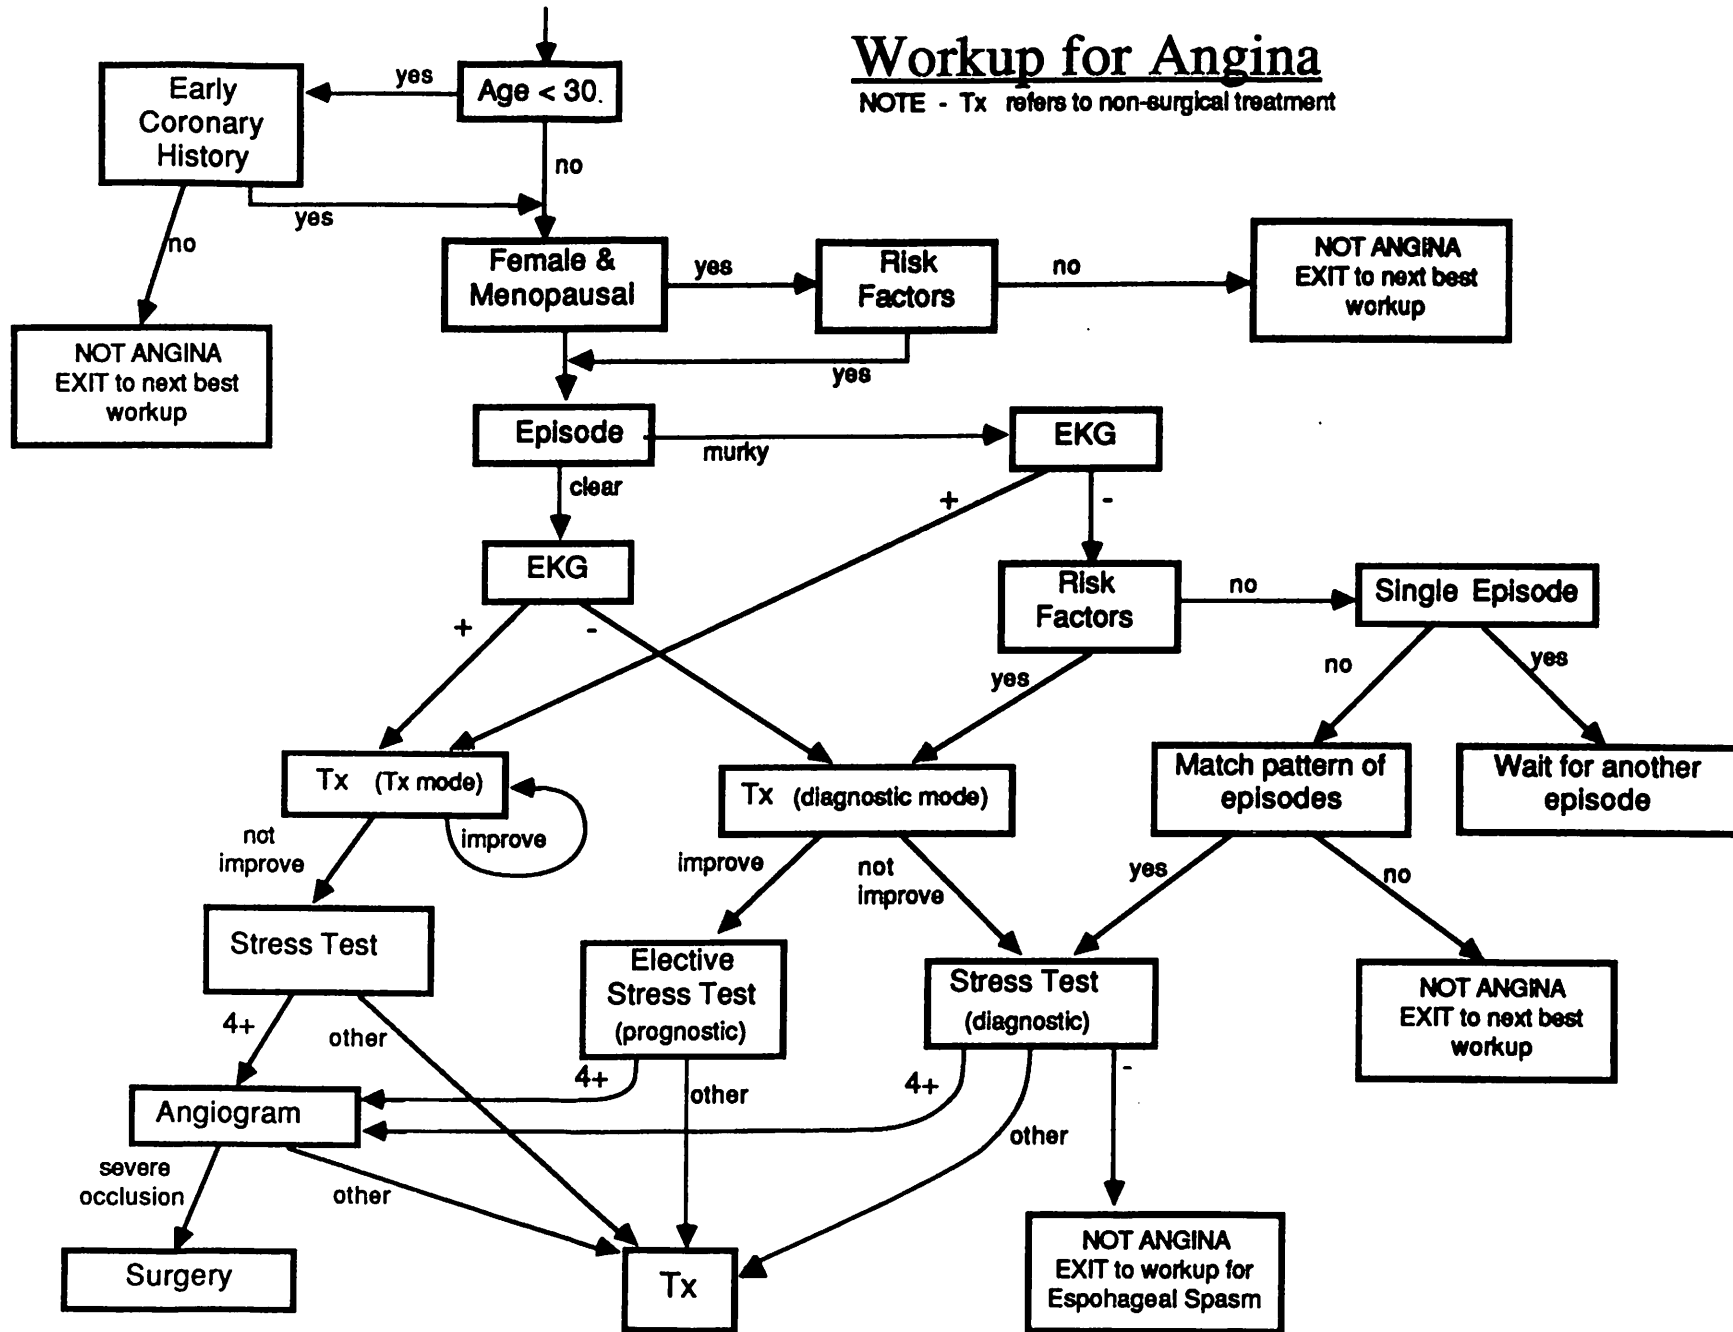
MUM diagnoses chest pain and abdominal pain. This includes taking a history, asking for physical findings, ordering tests, and prescribing trial therapy. Physicians call a diagnostic sequence of questions and tests a *workup*. MUM's *primary* goal is to generate workups for chest and abdominal diseases that include, in the correct order, all necessary questions and tests and none that are superfluous. Since we built MUM to study the management of uncertainty, the goal of correct diagnosis is secondary to generating the correct workup. We were influenced by a distinction physicians make between *retrospective* diagnosis, in which all evidence is known in advance and the goal is to make a correct diagnosis, and *prospective diagnosis*, which emphasizes the workup and proper management of the patient, even under uncertainty about his or her condition. MUM is definitely prospective. Figure 1 illustrates part of the workup for coronary artery disease. Clearly, we could build a system that follows this and other stored workups, but the point of the research is to be able to reason about the features of evidence, and the uncertainty in partially-developed diagnoses, to decide which questions to ask next. If MUM does this properly then its questioning will correspond with a standard workup, or at least be a reasonable alternative workup.

1.2 Managing Uncertainty and Control

MUM is based on the idea that managing uncertainty and controlling a complex knowledge system are manifestations of a single task, namely, acquiring evidence and using it to solve problems. There would be little basis for variation in problem-solving strategies if all evidence was equally costly, reliable, available, and pertinent; but if available and attainable evidence is differentiated along these and other dimensions, then problem-solving can be guided by the ideal of maximum evidence for minimum cost. For example, here is a strategy for focusing attention on available evidence:

Workup for Angina

NOTE - Tx refers to non-surgical treatment



CONTEXT: to minimize cost

CONDITIONS: $test_1$ and $test_2$ are pertinent, and
 $test_1$ is potentially-confirming, and
 $test_2$ is potentially-supporting, and
 $cost(test_1) \gg cost(test_2)$

ACTIONS: begin
do $test_2$
if supporting then do $test_1$
 else do not do $test_1$
end

That is, given cheap, weak evidence and expensive, strong evidence, get the weak evidence first and don't incur the cost of the strong evidence unless the weak evidence lends support. The rule serves to manage the uncertainty associated with the weak evidence – it says seek strong corroboration only if the weak evidence is positive. It also uses features of evidence such as cost and reliability to control the acquisition of evidence; for example, it explains why an angiogram (an expensive, risky, and excruciating test) is done only after a stress test in Figure 1. We distinguish these functions – managing uncertainty and control – only because uncertainty and control have, with a few exceptions noted below, been viewed as different topics. In fact, if control decisions are based on features of evidence, then control and managing uncertainty are the same thing. This is the principle that motivates the design of MUM discussed in Section 2.

1.3 Related Work

The close association between control and managing uncertainty has been apparent in the literature on sophisticated control for several years¹ but is largely absent from the AI literature on reasoning under uncertainty. Three important results have emerged from research on control: First, complex and uncertain problems have to be solved *opportunistically* and *asynchronously* – working on subproblems in an order dictated by the availability and quality of evidence (Hayes-Roth and Lesser, 1977). Second, since control tends to be accomplished by *local* decisions about focus of attention, the behavior of complex knowledge systems sometimes lacks global coherence. Coherence can be achieved by *planning* sequences of actions instead of selecting individual actions by local criteria². Third, programs are impossible to understand if the factors that affect control decisions are *implicit*. For example, the focus of attention in Hearsay-II was difficult to follow because it depended on many numerical parameters calculated from data and combined by empirical functions with “tuning” parameters (Hayes-Roth and Lesser, 1977). A better approach is to explicitly state and reason about the implicit factors, called *control parameters* (Wesley, 1983), that the numbers represent (Davis, 1985; Clancey, 1983). If the

¹For example, the classic paper by Erman, Hayes-Roth, Lesser, and Reddy (1980) is called “The Hearsay-II speech understanding system: Integrating knowledge to resolve uncertainty.”

²Personal communication, Victor Lesser.

control parameters are features of evidence and uncertainty, then control strategies can be developed to manage uncertainty.

This last point colors our reading of the AI literature on reasoning under uncertainty. Much of it is concerned with the mathematics of combining evidence, the calculation of *degrees of belief* in hypotheses. (A representative sample includes Shortliffe and Buchanan, 1975; Duda, Hart, and Nilsson, 1976; Zadeh, 1975; Shafer, 1976. See Cohen and Gruber, 1985; and Bonissone, 1985, for literature reviews, including nonnumeric approaches to uncertainty; and Szolovits and Pauker, 1978 for a discussion of uncertainty in medicine.) Degrees of belief *can* serve as control parameters, but it is necessary to maintain a distinction between combining evidence and control. Otherwise, degrees of belief (and the functions that combine them) have to be "tuned" not only to find the most likely answer but also to focus attention in a reasonable way. Inevitably they become ambiguous summaries of implicit control parameters. For example, MYCIN's certainty factors contained probabilistic and salience information, an indirect result of using them to focus attention (Buchanan and Shortliffe, 1985).

Another important reason to maintain the distinction between combining evidence and control is that combining evidence is only a part of the problem of reasoning under uncertainty. Other aspects include formulating decisions, assessing the need for more evidence, planning how to get it, deciding whether it is worth the cost and, if it isn't, hedging against residual uncertainty. In MUM we address the problem of combining uncertainty in the context of these other tasks.

2. An Architecture for Managing Uncertainty

Managing uncertainty in MUM requires many kinds of knowledge, discussed in this section. Anticipating section 2.3, on control, it may be useful to think of data moving bottom-up through Figure 2 as it triggers hypotheses and is requested by MUM's planner.

2.1 Types of Knowledge

Data, Evidence, and Interpretation Functions. Evidence is abstracted from data through interpretation functions. All data about a patient are stored in frames that describe personal history, family history, tests, history of episodes, and other data. Interpretation functions map data to evidence; for example, information that a patient smokes 3 packs of cigarettes a day is abstracted to the evidence *heavy-smoker* by an interpretation function that maps data about smoking habits to one of (*non-smoker light-smoker moderate-smoker heavy-smoker*) Interpretation functions are often graphs called *belief curves* that relate ranges of a continuous data variable to belief in evidence. Figure 3 shows a belief curve relating the duration of chest pain to the evidence *classic-anginal-pain*. Belief curves and other interpretation functions are acquired from an expert. They provide the same functionality as fuzzy predicates (Zadeh, 1975), and generalize Clancey's view of data abstraction as categorical (Clancey, 1983).

Features of Evidence. Evidence may be characterized by its cost, reliability, and roles. The cost of evidence reflects monetary cost as well as discomfort and risk to the patient (later

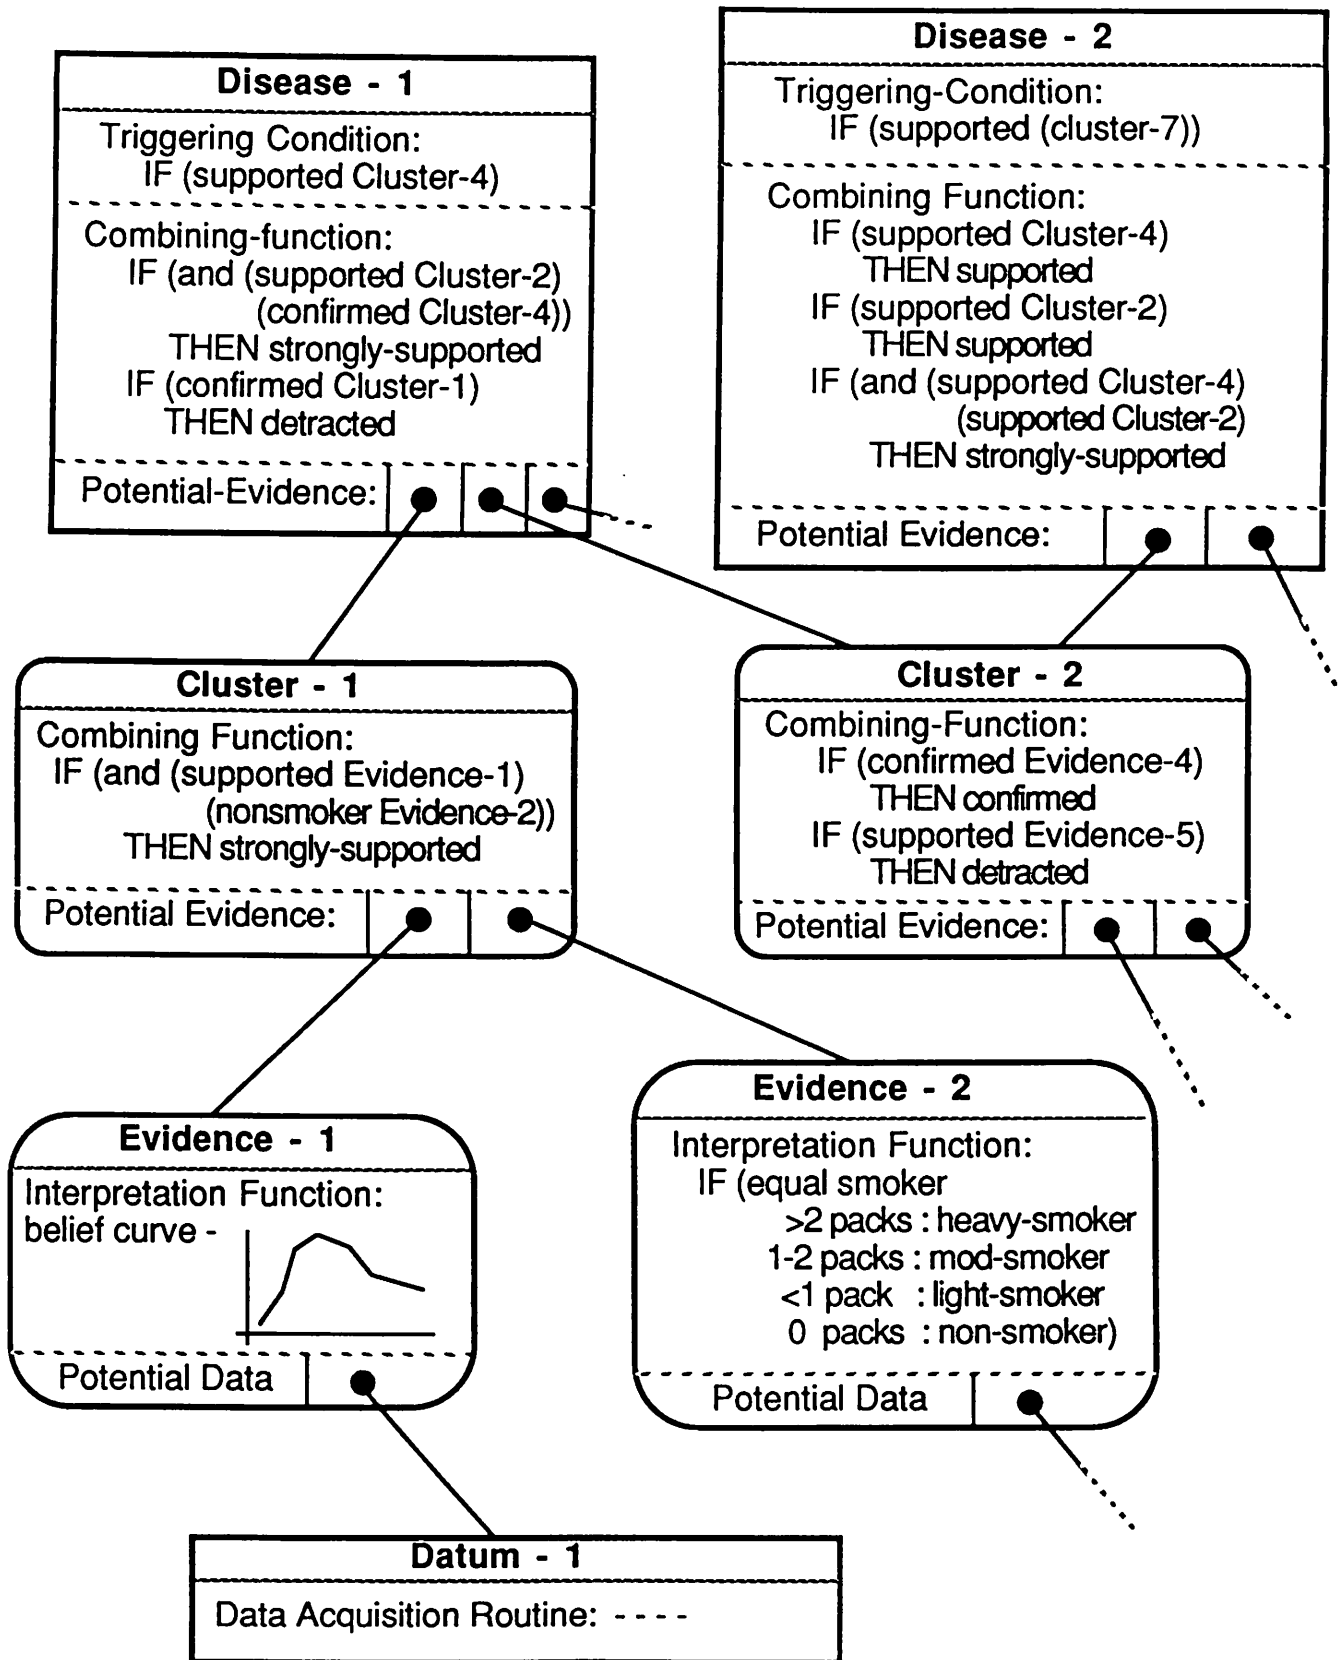



Figure 2: Knowledge Structures in MUM

Figure 3

A belief curve plotting the datum "Duration of Pain in Minutes"
vs. belief in the evidence "Classic-Anginal-Pain"



versions of MUM will separate these and other determinants of cost). Reliability refers to several factors, including false-positive and miss rates of tests, and also the belief in evidence derived from belief curves (e.g., is *classic-anginal-pain* at least *supported* by data about the pain duration?) The most important feature of evidence is the *roles* it can play with respect to evaluating hypotheses. MUM recognizes five roles, two of which are symmetric pairs:

Potentially-confirming and potentially-disconfirming. If evidence plays a potentially-confirming role with respect to a hypothesis, then acquiring it *might* confirm the hypothesis, though not all potentially-confirming evidence will, in actuality, confirm. For example, an EKG confirms the hypothesis of angina only if "positive" (i.e., shows ischemic changes.) Once confirmed (or disconfirmed), a hypothesis requires no further evidence, though a diagnostician may continue working to disconfirm other hypotheses, especially if they are dangerous.

Potentially-supporting and potentially-detracting. Like *potentially-confirming* and *potentially-disconfirming*, but not conclusive. However, combinations of supporting or detracting evidence may be confirming and disconfirming, respectively (see "Combining Functions," below). The combination referred to as cluster-2 (Fig. 2) is *potentially-supporting* with respect to disease-2; cluster-1 is *potentially-detracting* with respect to disease-1.

Trigger. A piece evidence plays the *triggering* role with respect to a hypothesis if its presence focuses attention on the hypothesis, or "brings the hypothesis to mind," or, in MUM, adds the hypothesis to a list of potential diagnoses. Cluster-4, if it is *supported* triggers disease-1 (Fig. 2). This role of evidence is found in virtually all medical expert systems.

Modifying. Some evidence does not support or detract from a hypothesis so much as it alters the way diagnosis proceeds. For example, risk factors for coronary artery disease (e.g., hypertension, elevated cholesterol) play a *modifying* role with respect to the hypothesis of angina since diagnosis will proceed aggressively if they are present and less aggressively otherwise.

These are the only roles currently used in MUM; others are contemplated. Note that evidence can play multiple roles with respect to any hypothesis; for example, risk factors are both *potentially-supporting* and *modifying* with respect to angina; and most *triggers* are individually or in combination with other evidence at least *potentially-supporting* (e.g., note the roles cluster-4 plays with respect to disease-1 in Fig. 2). Also, one piece of evidence can play different roles with respect to several hypotheses (illustrated by the roles cluster-2 plays with respect to disease-1 and disease-2 in Fig. 2). Finally, note that some evidence potentially plays two symmetric roles, while some is "asymmetric"; for example, a stress test will either support coronary artery disease or detract from it, while an EKG supports angina if it is positive and is useless otherwise. That is, EKG plays a *potentially-supporting* role only.

Clusters. Physicians often see collections of evidence that play particular roles in diagnosis; for example, shortness of breath that comes on suddenly but is unrelated to exercise or other inciting factors *triggers* the diagnosis of pulmonary embolism. Just as evidence has roles with

respect to clusters, so clusters have roles with respect to diseases, and these roles need not be supporting; for example, the cluster (*patient-age < 30 and no-family-history-of-coronary-events*) plays a *potentially-disconfirming* role with respect to all coronary diagnoses of chest pain. Instead of saying that the available evidence is a poor match to coronary diagnoses, we can say the evidence is a good match to a cluster that potentially detracts from or disconfirms coronary diagnoses.

Combining Functions. Every cluster includes a function, specified by the expert, that combines the available evidence for the cluster and returns a value for the cluster given evidence. The values returned by combining functions are just “realizations” of potential roles of evidence. For example, the value returned by the combining function of a cluster supported by *potentially-confirming* evidence could be *confirmed*. The value for a cluster with several pieces of *potentially-detracting* evidence might be *disconfirmed*, or perhaps *detracted*. Combining functions are further discussed below.

Diseases. A disease is technically a cluster. It is a collection of clusters, each of which plays an evidential role in diagnosis and is combined by combining functions with other clusters. Thus diseases reside at the top of a hierarchy of clusters (as shown in Fig. 2), each of which has its own combining function and specifications of the roles played by the clusters below it.

Strategic Knowledge. We characterize strategic knowledge as heuristics for deciding which triggered disease hypotheses to focus on, and how to go about selecting actions to gather evidence pertinent to these hypotheses. These heuristics have the same contingent nature as Davis’ meta-rules (Davis, 1985) and control rules in Neomycin (Clancey, 1985). We represent strategies as rules that include:

- *conditions* for selection of the strategy;
- a *focus policy* which guides the choice of a subset of the triggered disease hypotheses to focus on;
- *planning criteria* which guide the selection of actions to gather evidence and treat diseases currently in the focus.

Examples of focus policies are *plausibility* (choose hypotheses based on their degree of support); *criticality* (focus on hypotheses that, if true, would require immediate action); and *differential* (focus on hypotheses that offer alternate explanations for the symptoms). Examples of planning criteria are *cost* (prefer evidence that is easy to obtain, and inexpensive on some cost metric); *roles* (prefer potentially-confirming over potentially-supporting); and *diagnosticity*, meaning that a given result has the potential to increase the belief in one hypothesis and decrease belief in the other, as indicated by belief curves.

2.2 Combining Evidence and Propagating Belief

MUM combines evidence with local combining functions, as shown in Figure 2. Typically, knowledge systems require three functions to combine evidence and propagate belief. These are

illustrated in the context of two inference rules:

$$R1: (A \text{ AND } B) \rightarrow C$$

$$R2: (D \text{ AND } E) \rightarrow C$$

One function calculates the degree of belief (*dob*) in a conjunction from degrees of belief in the conjuncts:

$$\text{dob}(\text{AND } A \text{ B}) = F_1(\text{dob}(A)\text{dob}(B))$$

The second function calculates the degree of belief in a conclusion from

a) the degree of belief in its premise (computed by F_1)

b) the “conditional” degree of belief in the conclusion given the premise;

often called the degree of belief in the inference rule:

$$\text{dob}(C_{R1}) = F_2(\text{dob}(\text{AND } A \text{ B}), \text{dob}(C|(\text{AND } A \text{ B})))$$

The third increases the degree of belief in a conclusion when it is derived by independent inferences:

$$\text{dob}(C_{R1\&R2}) = F_3(\text{dob}(C_{R1}), \text{dob}(C_{R2}))$$

In MUM, these three kinds of combining are maintained, but with two important differences. First, there are no *global* functions corresponding to F_1 , F_2 , and F_3 ; all combining is done by functions local to clusters. Second, instead of the usual numeric degrees of belief, MUM has seven levels of belief: *disconfirmed*, *strongly-detracted*, *detracted*, *unknown*, *supported*, *strongly-supported*, *confirmed*. These are just “realizations” of the roles of evidence described earlier.

Combining evidence and propagating belief in MUM is illustrated in Figure 2. Each cluster, including diseases, has its own local combining function, specified by an expert. For example, cluster-1 is *strongly-supported* if the data *support* evidence-1 and if the data on a patient’s smoking habits support evidence that he or she is a nonsmoker. This is a conjunction of evidence of the kind calculated by F_1 , above. Another is found in the combining function for disease-1. If cluster-2 and cluster-4 are both *confirmed*, then disease-1 is *strongly-supported*. This illustrates the kind of combining for which F_2 , above, is required: even when the evidence for a disease is itself certain, the conditional belief in the disease given the evidence may not be certain. Disease-2 also contains a conjunctive rule, but the entire combining function illustrates the corroborative situation for which F_3 is needed. In this case, cluster-4 and cluster-2 *individually* play *potentially-supporting* roles, and taken together increase the level of belief in disease-2 to *strongly-supporting*.

Local combining functions have many advantages. Foremost is the ease with which an expert can specify *precisely* how the level of belief in a cluster depends on the levels of belief in the evidence for that cluster. Control of combining evidence is not relinquished to an algorithm, but is acquired from the expert as part of his or her expertise. Since local combining functions are

specific to clusters, they can be changed independently. And since the values passed between them in MUM are few, it is easy to trace back the derivation of a level of belief and pinpoint a faulty local combining function. The prospect of having to acquire many functions seems daunting, but we have found it easy and intuitive, and much easier to explain than a global numeric method.

2.3 Control of Diagnosis in MUM

MUM's basic control strategy involves three components:

User Interface: uses data description frames in the knowledge base to ask questions and create patient data frames for the results;

Matcher: uses the interpretation and combining functions to record the effect incoming data has on the belief states for clusters and disease frames, and triggers new hypotheses as appropriate;

Planner: uses strategic control rules from the knowledge base to guide the selection of focus and the planning process.

The planner controls the user interface and the matcher by requesting their services as described below.

Basic Control. The planner follows a basic control loop within which it interprets strategic control rules. It is implemented in a blackboard system with knowledge sources specified in the same syntax as strategic control rules. This facilitates experimental modifications. The design of the blackboard system was influenced by Hayes-Roth (1985), and shares the emphasis on explicit solution to the control problem. We first describe the basic control loop, then strategies and their selection.

The basic control loop is initiated with the choice of a *strategic phase*. All strategic phases but one include a *focus policy* that directs MUM's attention to a subset of candidate hypotheses. This is followed by the selection of short-term plans to gather evidence and select treatment pertinent to these hypotheses (the rule in Section 1.2 represents such a plan). Since the effort of developing lengthy plans may well be wasted in a domain permeated with uncertainty, we constrain plans to single actions or sequences of two actions where the applicability of the second depends on the outcome of the first. Several short-range plans may be generated and executed.

Carrying out plans typically consists of invoking the user interface to request some information, updating the status of the diseases with the matcher, and conditional continuation of the plan. When no short-term plans remain, the system iterates the basic control loop to determine if a new strategic phase is appropriate, update the focus, and generate new short-term plans.

Strategic Control. We represent MUM's overall strategy as an ordered set of rule-like strategic phases, shown in Figure 4. Each phase has conditions that activate it. Once activated, a phase controls MUM's focus of attention and the choice of actions pertaining to the hypotheses in this focus.

Strategic Phase:	Get General Picture.
Conditions:	No candidate hypotheses.
Focus Policy:	None.
Planning Criteria:	Evidence must play trigger role; prefer low cost on all cost metrics.
Strategic Phase:	Initial Assessment for Triggered Hypotheses.
Conditions:	One or more hypotheses are triggered.
Focus Policy:	Focus on triggered hypotheses.
Planning Criteria:	Must be low on all cost metrics; prefer stronger roles.
Strategic Phase:	Deal With Critical Possibilities
Conditions:	There are critical hypotheses which have not been confirmed, disconfirmed or strongly detracted, and if they are detracted, no other hypothesis is confirmed.
Focus-Policy:	Criticality.
Planning Criteria:	Rule Out if possible, else gather support. Utility of evidence. Low cost first; as needed let discomfort and monetary cost increase.
Strategic Phase:	Discriminate Strongest Hypotheses
Conditions:	More than one hypothesis is supported.
Focus-Policy:	Plausibility.
Planning Criteria:	Diagnosticity, Low cost first. Utility of evidence. Substitute high cost confirmation for one hypothesis with lower cost disconfirmation for the other.

Figure 4: Four Strategic Phases in MUM's Diagnosis

The phase **Get General Picture** is invoked when the system is started, and may also be used if all previously considered hypotheses are ruled out. It has no focus policy because no hypotheses are active when it is invoked. It directs the planner to ask for evidence that plays the **potential-trigger** role for one or more hypotheses, pursuing the lowest-cost evidence first. The cluster **initial-consultation** (consisting of age, sex, and primary complaint) meets the criteria of potentially triggering many hypotheses and costing little. The initial consultation usually triggers some hypotheses, which result in a new strategic phase being selected. If no hypotheses were triggered, the planner asks for potential-triggers of higher cost.

The **Initial Assessment for Triggered Hypotheses** phase is invoked when new hypotheses are triggered. Since the conditions of the other strategic phases depend somewhat on the level of belief in candidate hypotheses, this phase gathers preliminary evidence for the hypotheses. The focus is on the triggered hypotheses, so only evidence playing some role relative to these hypotheses is considered by the planner. This phase directs the planner to gather low-cost evidence for the hypotheses. For example, MUM asks about aspects of the patient's episode (the event which is the primary complaint) which bear on the triggered hypothesis, and about risk factors.

As soon as the easy questions for triggered hypotheses have been asked, MUM decides between the next two phases based its belief in the hypotheses and whether any of the hypotheses are *critical*, that is, require immediate treatment if supported. Critical hypotheses are dealt with first.

The **Deal With Critical Hypotheses** phase places all candidate critical hypotheses in MUM's focus. The short range planner is then directed to attempt to rule out these hypotheses. It begins with potentially-disconfirming or potentially-detracting evidence. If it fails to find any, then it looks for potentially-supporting evidence. It will not seek evidence that plays a lesser potential role than evidence it already has. For example, it will not seek potentially-supporting evidence for a hypothesis that is already strongly supported, but rather focuses on potentially-confirming evidence. The planner will focus on low-cost evidence first, but it is not prohibited from pursuing high-cost evidence as it was in the previous phase.

If the focus of attention is not captured by critical hypotheses, it is dictated by plausibility. The strategic phase **Discriminate Strongest Hypotheses** discriminates competing alternatives with as little cost to the patient as possible. As before, the potential roles of evidence are used to decide whether it is worth acquiring.

Currently MUM stops work when a hypothesis is confirmed and no critical hypotheses remain in its focus. We are implementing the next strategic phases, prognosis and treatment. Both provide evidence of diagnostic significance; for example, MUM may begin treatment for angina if it is strongly supported, rather than incur the cost of absolute confirmation. If the treatment relieves the symptoms, then it is additional evidence for the diagnosis. If not, it is evidence that detracts from the diagnosis and may support others. Since treatment provides evidence, we represent treatments as clusters, exactly the same way as we represent tests such as angiography.

The emphasis in MUM is on asking the right questions in the right order without superfluous questions. MUM's control knowledge is not yet sophisticated enough to satisfy all these criteria. It asks questions in a reasonable order, but it sometimes focuses on the wrong disease. Since MUM is a nascent system, this does not yet concern us. We believe the system is successful in

providing a framework for exploring management of uncertainty by sophisticated control, that is, by making control decisions based on the roles, costs and other characteristics of evidence; and the criticality of diseases; and the credibility of diagnoses.

3. Conclusions

MUM manages uncertainty by reasoning about evidence and its current state of belief in hypotheses. Its goal is to generate appropriate workups for chest and abdominal pain, that is, to ask the right questions in the right order without unnecessary questions. To the extent it succeeds, it demonstrates its ability to manage uncertainty, to select the appropriate action given uncertainty. We have said this is a control task. Indeed, much of MUM's architecture is devoted to explicit, evidence-based control.

Much work remains to be done. Currently, MUM resembles a programming environment more than a medical expert system. We will be devoting ourselves to building up its clusters and control rules.

Although MUM was designed for medical problems and is discussed in that context, we believe the approach to uncertainty and control it engenders is general to classification problem solvers, as well as to other systems responsible for the management of uncertainty. An empty version of MUM called MU is being developed and will be tested in other domains.

REFERENCES

- [1] Bonisonne, P. 1985. Reasoning with uncertainty in expert systems. *International Journal of Man-Machine Studies*, 22:3.
- [2] Clancey, W. 1983. The advantages of abstract control knowledge in expert systems design. In *Proceedings of the Third National Conference on Artificial Intelligence*.
- [3] Cohen, P. and Gruber, T. 1985. Reasoning about uncertainty: A knowledge representation perspective. Pergamon Infotech State of the Art Report.
- [4] Davis, R. 1985. Interactive transfer of expertise. In *Rule-based expert systems*, B. Buchanan and E. Shortliffe, (Eds.) Addison-Wesley.
- [5] Duda, R.O. Hart, P.E. and Nilsson, N. 1976. Subjective Bayesian methods for rule-based inference systems. Technical note 124, AI Center, SRI International, Menlo Park, CA.
- [6] Hayes-Roth, B. 1985. A blackboard architecture for control. *Artificial Intelligence*, Vol. 26, pp. 251-321.
- [7] Hayes-Roth, F. and Lesser, V. 1977. Focus of attention in the Hearsay-II speech understanding system. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*.
- [8] Patel, V. and Groen, G. 1986. Knowledge based solution strategies in medical reasoning. *Cognitive Science*, Vol. 10, pp. 91-116.
- [9] Shafer, G. 1976. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, New Jersey.
- [10] Shortliffe, E. and Buchanan, B. 1975. A model of inexact reasoning in medicine. *Mathematical Biosciences*, Vol. 23, pp. 351-379.
- [11] Szolovits, P. and Pauker, S.G. 1978. Categorical and Probabilistic Reasoning in Medical Diagnosis. *Artificial Intelligence*, V.11, pp. 115-144
- [12] Wesley, L.P. 1983. Reasoning about control: the investigation of an evidential approach. *Proceedings IJCAI-83*, pp. 203-206.
- [13] Zadeh, L.A. 1975. Fuzzy logic and approximate reasoning, *Synthese*, Vol. 30, pp. 407-428.