

SRI International



REASONING ABOUT CONTROL: AN EVIDENTIAL APPROACH

Technical Note 324

July 30, 1984

By: Leonard P. Wesley Computer Scientist*
John D. Lowrance, Assistant Director AIC
Thomas D. Garvey, Staff Scientist

Artificial Intelligence Center
SRI International
Menlo Park, California 94025

The work reported here was supported in part by the Defense Advanced Research Projects Agency (DARPA) under Contract No. N00014-81-C-0115, (DARPA) under Contract No. N00014-82-K-0464, and by the Naval Electronic Systems Command under Contract No. N00039-83-K-0656 (SRI Project No. 6486). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of SRI International, DARPA, the Navy, or the U.S. government.

* The author is currently enrolled in the department of Computer and Information Science (COINS), University of Massachusetts, Amherst, Massachusetts 01003.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 30 JUL 1984		2. REPORT TYPE		3. DATES COVERED 00-07-1984 to 00-07-1984	
4. TITLE AND SUBTITLE Reasoning About Control: An Evidential Approach				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SRI International, 333 Ravenswood Avenue, Menlo Park, CA, 94025				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 42	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

ABSTRACT

Expert systems that operate in complex domains are continually confronted with the problem of deciding what to do next. Being able to reach a decision requires, in part, having the capacity to "reason" about a set of alternative actions. It has been argued that expert systems must reason from *evidential* information—i.e., *uncertain*, *incomplete*, and occasionally *inaccurate* information [LOW82a]. As a consequence, a model for reasoning about control must be capable of performing several tasks: to combine the evidential information that is generically distinct and from disparate sources; to overcome minor inaccuracies in the evidential information that is needed to reach a decision; to reason about what additional evidential information is required; to explain the actions taken (based on such information) by the system. These are a few of the formidable control problems that remain largely unsolved [BAR82]. If expert systems are to improve their performance significantly, they must utilize increasingly sophisticated and general models for dealing with the evidential information required for reasoning about their behavior. To this end we present an alternative evidentially-based approach to reasoning about control that has several advantages over existing techniques. It enables us to reason from limited and imperfect information; to partition bodies of meta- and domain-knowledge into modular components; and to order potential actions flexibly by allowing any number of constraints (i.e., control strategies) to be imposed over a set of alternative actions. Furthermore, because it can be used for reasoning about the expenditure of additional resources to obtain the evidential information needed as a basis for choosing among alternatives, this approach can be employed recursively.

INTRODUCTION

In complex domains, there are typically several alternative actions a system might take to complete its tasks. In some cases, the best alternative is clear or the choices do not warrant extensive analysis. At times the consequences of pursuing some action justify expending the effort to analyze the pros and cons of choosing a particular alternative. This is especially true when the effect of some choice is potentially counterproductive and irrevocable, or might result in the inefficient use of a system's limited resources.

When such an analysis is warranted, a system must be capable of "reasoning" about its actions from *uncertain*, *incomplete*, and occasionally *inaccurate* information called *evidential* information [LOW82a]. For instance, the expected outcome of pursuing a particular action is one of many factors that can influence the decision to choose that action. However, situations may arise in which *uncertainty* exists about the consequences of pursuing *any alternative*, particularly when uncontrollable or unpredictable events may intervene. Indeed, a system will frequently "find" that it is uncertain about much of the information required for reaching a decision.

There are several reasons a system must reason from *incomplete* information. One explanation is that resource limitations might prevent the system from gathering all relevant facts, thus possibly forcing choices to be made on the basis of limited (i.e., incomplete) evidence. A second reason is that some information can be lost during its translation or its abstraction into alternative representations, and so it is less complete. This becomes readily apparent in such task domains as general-purpose computer vision, and speech and natural-language understanding. A third reason is that information-gathering processes are far from perfect; they are consequently limited in their ability to extract all of the available critical information from their environment.

Finally, it must be anticipated that evidential information will be occasionally *inaccurate* because, among other reasons, the sources of the information (procedural and declarative) are also imperfect.

It can thus become a nontrivial problem to make choices on the basis of evidential information. Yet we, as well as the systems built by us, must do exactly that in making all kinds of important decisions in real-world situations.

In support of our position, consider the following problem. A physician does not know if his patient's soreness is caused by a viral or bacterial infection. If a bacterial infection is the cause, penicillin might be prescribed; otherwise rest and liquids might be the appropriate treatment for a viral infection. Failure to treat a bacterial infection could complicate the patient's situation or, in some cases, result in his death. However, he might be allergic to certain antibiotics or the bacteria could possibly be penicillin-resistant.

A culture taken from the patient can be tested to determine whether bacteria are present and, if so, what type. However, tests are not always accurate—a culture can die or be contaminated before reaching the laboratory; inappropriate tests might be performed or bad laboratory procedures employed. Furthermore, the mere presence of bacteria is not conclusive evidence that the patient's illness is due to a bacterial infection.

What should the physician do? Let us consider a subset of the possible actions, from which one must be chosen:

- Take no culture and treat the illness as a viral infection.
- Take no culture and prescribe penicillin.
- Take a culture and request some tests for bacteria; if results are positive, prescribe penicillin.
- Take a culture and request some viral tests; if results are positive, prescribe rest and liquids.
- Take a culture and request some tests for both bacteria and viruses, then prescribe penicillin if tests for bacteria are positive or rest and liquids if viral tests are positive.
- Take a culture and request extensive tests for bacteria and viruses, then prescribe penicillin if bacteria tests are positive; otherwise, if viral tests are positive, prescribe rest and liquids.

The consequences of taking any of the above actions cannot be known with complete certainty. However, it might appear that, for any given action, some consequences are more likely or acceptable than others. Therefore, a decision must be reached, to at least some extent, on the basis of partial beliefs about the possible and acceptable outcomes associated with taking each action.

Although knowing the exact cause of the patient's condition would help significantly in deciding what actions are necessary, such complete information is usually not available. For example, if only the morphology of the bacteria is known, then their identity must be subsumed within the set of bacteria that could possibly conform to the observed dimensions. On the basis of this shape information, the doctor might only be willing to express his belief that the presence of some types of bacteria is more likely than others. Consequently, he might prefer to suggest that some actions are more likely to be effective than others in ultimately curing his patient. This limited information about shape is insufficient for complete and precise identification of the bacteria. In fact, it is because the physician is partially ignorant about some of the possibilities that he might take certain actions to become more knowledgeable (e.g., request a Gram stain test).

Another factor compounding the doctor's problem is that the information needed to choose his next course of action may be inaccurate. If the morphological information is in gross error, subsequent efforts may result in misidentification of the bacteria and thus possibly a misdiagnosis. Such errors, if and when found, have to be discarded or outweighed by redundant information from other sources, such as additional tests. Alternatively, if the shape information is generally correct, the diagnosis is more likely to be valid and the appropriate treatment administered.

The doctor's decision to perform certain actions is no more effective than the certainty, completeness, and accuracy of the information needed for choosing among alternatives. Furthermore, this scenario illustrates that the luxury of having perfect information is at best rare in complex domains, and that the real world requires expert systems to be capable of reasoning about and explaining their behavior in accordance with a potentially large and diverse body of evidential information.

This fact requires that a model for reasoning about control have the following capabilities: (1)combine generically distinct evidential information that has been obtained from disparate sources; (2)correct for minor errors in bodies of evidence; (3)be flexible in its ability to order its actions by allowing any number of constraints (i.e., control strategies) to be imposed upon a set of alternatives; (4)be employed recursively in that it may be used to reason about taking further actions to obtain the evidential information needed for making choices; (5)explain its decisions more accurately and completely.

ACTIONS, CONTROL-FEATURE SPACES, CONTROL ENVIRONMENT,
AND CONTROL KNOWLEDGE SOURCES (CKS)

We consider an action to be one of the following:

- (1) The invocation of a parameterized process, such as a knowledge source (KS); a KS in this context is a procedure or sensor that makes observations about the environment.
- (2) The invocation of processes, such as control knowledge sources (CKS), to gather additional information needed to choose from a set of alternative actions.
- (3) The activation of a parameterized effector that may change the state of the environment. (This is a special case of No. 1 above.)

To choose a course of action, a system must be capable of making observations about features in the environment. Just as a tumorous cell exhibits certain morphological features that can help distinguish it from healthy cells, so can actions exhibit certain "control features" that can help distinguish those alternatives that are more desirable than others. For example, fibrous histiocytomas are a particular type of tumorous cell that exhibits certain distinguishing morphological features—e.g., they characteristically exhibit a proliferation of fibroblasts admixed with multinucleated giant cells—see Figure 1. Healthy cells do not exhibit this shape characteristic.

In general, fibrous histiocytomas in superficial locations usually follow a benign clinical course. However, if the tumor develops in deep tissue, it tends to be malignant. But depth is a relative and sometimes quite subjective measure. What may be considered deep in thigh muscle tissue might not be so considered in forearm




Figure 1.

CHARACTERISTIC MORPHOLOGY OF FIBROUS HISTIOCYTOMA TUMORS

muscle tissue. Furthermore, for any particular biopsy it is not unusual for physicians to differ in their beliefs as to whether it *is* or *is not* deep. If the doctor elects not to operate and remove the tumor, there is a nonzero probability of the tumor's becoming malignant.

In this example the depth of the tumor can be considered a control feature that may help doctors to decide which course to pursue—whether to remove the tumor or leave it in the patient. A control feature, therefore, is any observable or quantifiable aspect of an environment and system that could possibly help to discern the appropriate action.

Examples of “control feature spaces” are goals/subgoals, and costs associated with invoking various KSs, the reliability of KSs, as well as, for instance, the superficiality of tissues. Each control feature space can be viewed as a set of at least one observable and quantifiable “control feature value.” For our superficiality control feature space, two possible control feature values are “deep” and “not deep”. In short, the extent to which a system is capable of choosing the appropriate action

depends on its ability to observe and make measurements of control features.

The set of all control feature spaces of potential interest constitutes a "control environment." In general, this includes any aspect of a system and its environment about which information may be obtained in order to realize a decision.

For our approach, the information required to reason about what to do next is obtained through a set of control knowledge sources.¹ Similar to KSs, CKSs provide partially processed, "control-related," evidential information that is based on the CKSs' observations of the control environment.

A FORMAL VIEW OF THE CONTROL PROBLEM

Let A , the set of mutually exclusive and exhaustive actions that a system can possibly take, be defined as

$$A = \{a_1, a_2, \dots, a_n\}.$$

Let F_1, F_2, \dots, F_m constitute the control feature spaces of interest.² Associated with each F_i , $1 \leq i \leq m$, is a set \mathcal{F}_i of possible feature values of F_i , where

$$\mathcal{F}_i = \{f_i^k \mid f_i^k \text{ is a possible feature value of } F_i \text{ where } 1 \leq k \leq |\mathcal{F}_i|\}.$$

For each $f_i^k \in \mathcal{F}_i$, it is possible to identify a subset of actions whereby each action is possibly the best one to take when f_i^k is observed. For instance, if measurements indicate that the goal of getting from point A to point B as fast as possible should be satisfied, then flying or driving, rather than walking, is possibly

¹ A CKS differs from a KS only in the scope of the environment and types of features it is expected to perceive.

² Because we are limited in our capacity to reason with large bodies of information, a premise of our research is that we deal with this dilemma by employing methods for producing a manageable set of alternative actions and control-related information. We suggest that analogous methods must be employed in expert systems.

the best action. However, if measurements indicate that the goal of getting from point A to point B by relatively economical means should be satisfied, then walking is possibly the best action.

We can now construct a frame of action, Θ_a :

$$\Theta_a = \{(a_j, f_1^k, f_2^{k'}, \dots, f_m^{k''}) \mid a_j \text{ is possibly the best action to take when } f_1^k, f_2^{k'}, \dots, \text{ and } f_m^{k''} \text{ are observed}\},$$


$$\Theta_a \subseteq A \times \mathcal{F}_1 \times \mathcal{F}_2 \times \dots \times \mathcal{F}_m.$$

Each \mathcal{F}_i can be thought of as inducing an equivalence relation over Θ_a that partitions Θ_a into $|\mathcal{F}_i|$ equivalence classes. Each $f_i^k \in \mathcal{F}_i$, therefore, is effectively an equivalence class whose members are those elements of Θ_a that exhibit feature value f_i^k . Therefore, $f_i^k \subseteq \Theta_a$ and, as a consequence,

$$\Theta_a = \bigcup_{f_i^k \in \mathcal{F}_i} f_i^k.$$

Let us illustrate this formalization in Figures 2 through 4 below. Consider the Venn diagram in Figure 2. In this figure, Θ_a is partitioned into four equivalence classes whereby, for instance, if f_1^1 is observed, then taking an action that is a member of f_1^1 is more appropriate than taking any action that is not a member of f_1^1 .

Figure 3 illustrates how the control feature space \mathcal{F}_3 may partition Θ_a differently from \mathcal{F}_1 . Figure 4 is an illustration of how \mathcal{F}_1 and \mathcal{F}_2 partition Θ_a and of the fact that there are some actions, for example (a_{54} , a_{55} , and a_{56}), that neither \mathcal{F}_1 nor \mathcal{F}_2 alone can completely distinguish. When just f_2^1 is observed, only a_{54} and a_{55} can be distinguished from a_{56} . Similarly, just observing f_1^3 allows only a_{55} and a_{56} to be distinguished from a_{54} . However, if f_1^3 and f_2^1 are observed



NOTE: a_i 's label grid cells; f_i 's label areas bounded by.

Figure 2.

A SAMPLE FRAME OF ACTION, $\Theta_A \subseteq A \times \mathcal{F}_1$. WHERE $\mathcal{F}_1 = \{f_1^1, f_1^2, f_1^3, f_1^4\}$.

simultaneously, a_{54} , a_{55} and a_{56} can each be distinguished. We see that, if a system is to choose among alternative actions, it must know the features and perhaps their values that must be observed in order to make a choice.

A Simple Example

In this section we show how this formalization might be applied to the previous medical problem. Consider the following subset of possible actions of A :

A) ACTIONS:

- a_1) Take no culture and treat the illness as a viral infection.
- a_2) Take no culture and prescribe penicillin.

a_1	a_2	...		f_2^4					
		f_2^3					f_2^1		
			a_{54}	a_{55}	a_{56}				
				f_2^2			f_2^5	...	a_n

Figure 3.

A SAMPLE FRAME OF ACTION, $\Theta_A \subseteq A \times \mathcal{F}_2$. WHERE $\mathcal{F}_2 = \{f_2^1, f_2^2, f_2^3, f_2^4, f_2^5\}$.

- a_3) Take a culture and request some tests for bacteria; if results are positive, prescribe penicillin.
- a_4) Take a culture and request some viral test; if results are positive, prescribe rest and liquids.
- a_5) Take a culture and request some tests for both bacteria and viruses, then prescribe penicillin if tests for bacteria are positive or rest and liquids if viral tests are positive.
- a_6) Take a culture and request extensive tests for bacteria and viruses, then prescribe penicillin if bacteria tests are positive; otherwise, if viral tests are positive, prescribe rest and liquids.

and the following control feature spaces,

f_2^2) The physician's goal/subgoal should be to make a diagnosis at moderate cost to the patient.

f_2^3) The physician's goal/subgoal should be to make a diagnosis regardless of cost to the patient.

F_3) TEST RELIABILITY

f_3^1) Tests for bacteria are very reliable.

f_3^2) Tests for bacteria are not very reliable.

F_4) RISK TO HUMAN LIFE

f_4^1) The risk of delayed treatment is relatively great.

f_4^2) The risk of delayed treatment is relatively moderate.

f_4^3) The risk of delayed treatment is relatively low.

Let us determine the extent to which our choice of control feature spaces is governed by intuitive factors. Consider F_1 , based on the physician's evaluation of the patient's situation; he might believe that the distinctions within a subset of alternative treatments are so subtle that far more tests than usual are required. He would want to consider only those actions that would allow such distinctions to be made and exclude others that cannot. Extensive tests to determine the correct treatment are therefore required.

Second consider F_4 . A situation might arise in which a patient appears to be so severely ill that delayed treatment could be fatal—that is, there is seemingly no time to take and analyze cultures. After assessing the situation, the physician might favor some action, such as a_1 or a_2 , that does not involve lengthy laboratory tests.

Given plausible arguments for our remaining choices of these control feature

spaces, the following frame of action can be constructed:

$$\Theta_a = \{(a_1, f_1^1, f_2^1, f_3^2, f_4^1) \\ (a_2, f_1^1, f_2^1, f_3^2, f_4^1) \\ (a_3, f_1^2, f_2^2, f_3^1, f_4^2) \\ (a_4, f_1^2, f_2^2, f_3^2, f_4^2) \\ (a_5, f_1^2, f_2^2, f_3^1, f_4^3) \\ (a_6, f_1^3, f_2^3, f_3^1, f_4^3)\}.$$

Our reason for indicating that actions a_1 and a_2 are possibly the most appropriate ones when the control feature values f_1^1, f_2^1, f_3^2 , and f_4^1 are observed is that, if the doctor believes

- That the diagnosis should be made with a minimum number of tests or none at all,
- That his goal/subgoal should be to keep the patient's costs as low as possible,
- That the tests to determine the presence, absence, or type of bacteria are not very reliable, and
- That delayed treatment might be fatal to the patient,

then a_1 and a_2 are consistent with these beliefs. Our reason for indicating that actions a_1 and a_2 are not possibly appropriate when the control feature value f_4^3 , for example, is observed is that, if the doctor believes that the risk of delayed treatment is relatively low, it is not appropriate to prescribe a treatment without performing some tests to be more certain about the cause of the problem. The reasons for our remaining choices are similarly motivated.

This example shows that there might be situations in which the doctor will not be able to choose among actions—for instance a_1 and a_2 when f_1^1, f_2^1, f_3^2 , and f_4^1 are observed. Without additional observations of control features that can distinguish a_1 from a_2 , there is no justification for choosing one over the other. Under these circumstances, an arbitrary selection must be made.

As the example stands, CKSs must convey their observations in terms of the individual f_i^k 's. Realistically, CKSs cannot always determine if, for example, the

risk to human life is great, moderate, or low. Occasionally the best observation that can be made is that the risk is *not* low (i.e., *possibly* moderate or great). A CKS, therefore, must be able to communicate in terms of disjunctions of possibilities—for instance, the proposition $f_4^1 \vee f_4^2$. Furthermore, at times a CKS may need to express its total ignorance about a proposition. For example, the risk assessment CKS can express total ignorance through the proposition $f_4^1 \vee f_4^2 \vee f_4^3$. It is important that a model for reasoning incorporate those propositions through which CKSs wish to convey their observations.

Realistically, we can expect a model to provide only a partial ordering over a set of alternative actions. This can be illustrated with our medical example. To the degree our observations indicate that tests for bacteria are not reliable (i.e., f_3^2 is true), we should favor any one of the actions $\{a_1, a_2, a_4\}$ over any of the actions $\{a_3, a_5, a_6\}$. If additional observations indicate that the risk of delayed treatment is relatively moderate, then we should favor action a_4 over any of the actions $\{a_1, a_2, a_3, a_5, a_6\}$. We see that a partial ordering of actions is produced as a consequence of the type and strength of observations made about the environment.

If CKSs were perfect, the control problem would take on a different form because the value of every feature could be determined—and thus the appropriate action more easily selected. However, as discussed, perfect information is typically not available. A system cannot always obtain the exact value for a subset of features, nor will it always know what subset of features is required for choosing among alternative actions. At best, a system will be able to refine and order a subset of alternatives only to the extent that it is capable of realizing what control features are necessary to discern the appropriate action and is able to reason from imperfect measurements of those features.

REASONING FROM EVIDENTIAL INFORMATION

Information used by expert systems to reason about complex environments is evidential—i.e., it is often uncertain, incomplete, and inaccurate. This follows from the fact that the world is perceived through a set of KSs that provide only partially processed sensory information. Furthermore, such information is inherently

evidential and not readily expressed in terms of simple truths and falsities. Just as the information needed to interpret our environment is evidential in nature, so is the information needed to decide what to do next. That is, information about such things as goals, consequences of actions, the power of KSs, knowledge about the environment, and so on are not easily expressed in terms of Boolean logic or probabilities. Rather, such information is evidence that tends to confirm or refute hypotheses about what control feature values have been observed, and these hypotheses in turn imply, directly or indirectly, what actions should be taken.

BODIES OF EVIDENCE AND MASS DISTRIBUTIONS

Up to this point we have stated that a CKS may communicate about its observations by providing a body of evidence that expresses its beliefs about propositions. In this section we discuss the form of such evidence.

A CKS conveys its observations in the form of a "mass distribution," which it derives from a body of evidence that tends to confirm or refute a subset of the propositions of interest. In an evidential approach, every CKS has a unit of "mass" that it may distribute, on the basis of its beliefs about what it has observed, among the various propositions of interest. Given its observations, a CKS might believe that a subset of these propositions is partially or completely true. Such beliefs can be conveyed by attributing a proportionate amount of the CKS's unit mass directly to the truthfulness of those propositions. Conversely, a CKS can attribute a portion of its mass directly to the negation of a subset of propositions if it believes that it is partially or completely false.

Bayesian and Mass Distributions

A mass distribution can be viewed as a generalized Bayesian distribution of belief over a set of propositions. A Bayesian distribution assigns a unit of belief over a set of mutually exclusive and exhaustive propositions, as designated by the mapping m :

$$m : \Theta_a \mapsto [0, 1], \quad \text{where } \sum_{p \in \Theta_a} m(p) = 1.$$

The probability of any proposition, for instance $B \subseteq \Theta_a$, is the sum of the belief attributed to propositions that imply B or one minus the sum of the belief attributed to propositions that imply not B (i.e., $\neg B$):

$$\text{for all } B \subseteq \Theta_a, \quad \text{Prob}(B) = \sum_{p \in B} m(p).$$

It follows that

$$\text{Prob}(B) = 1 - \text{Prob}(\neg B).$$

However, a mass distribution M need not attribute belief to mutually exclusive and exhaustive propositions:

$$M : 2^{\Theta_a} \mapsto [0, 1], \quad \text{where } M(\emptyset) = 0 \quad \text{and}$$

$$\sum_{p \subseteq \Theta_a} M(p) = 1.$$

The sum of the mass attributed to propositions that imply B plus the sum of the mass attributed to propositions that imply $\neg B$ need not equal one, because some mass may have been assigned to propositions that imply neither (e.g., Θ_a).

Each body of evidence, therefore, induces an interval, called an "evidential interval," within which belief about a proposition must lie. An evidential interval is a subinterval of the real interval $[0,1]$. The lower and upper bounds of the evidential interval will be called support (*Spt*) and plausibility (*Pls*), respectively. The *Spt* represents the total mass that tends to support a proposition:

$$\text{Spt}(B) = \sum_{p \subseteq B} M(p).$$

The Pls represents the degree to which the mass fails to refute the proposition:

$$Pls(B) = 1 - Spt(\neg B) = 1 - \sum_{p \subseteq \neg B} M(p).$$

The degree to which the mass refutes a proposition is called its dubiety (Dbt):

$$Dbt(B) = Spt(\neg B) = 1 - Pls(B).$$

The degree to which no mass is attributed to a proposition or its negation is called ignorance (Igr):

$$Igr(B) = Pls(B) - Spt(B).$$

The interpretations of some evidential intervals are summarized below:

Completely true proposition $[1, 1]$;

Completely false proposition $[0, 0]$;

Completely ignorant about the proposition $[0, 1]$;

Tends to support the proposition $[Spt, 1]$, $0 < Spt < 1$;

Tends to refute the proposition $[0, Pls]$, $0 < Pls < 1$;

Tends to both support and refute $[Spt, Pls]$, $0 < Spt \leq Pls < 1$.

Combining Bodies of Evidence

Typically several distinct CKSs must be invoked to obtain the evidence required to decide upon a course of action. This follows from the fact that no single CKS is knowledgeable, in general, about all f_i^k 's $\in \mathcal{F}_i$'s. Any reasoning mechanism, therefore, must be capable of combining evidence from distinct CKSs. The combining rule that was developed by Dempster and which is central to our approach is applied to two bodies of evidence, in the form of mass distributions—e.g., M_1 and

M_2 —and produces a third mass distribution $M_{1\oplus 2}$ [DEM67, DEM68]:³

$$\text{For all } B_1, B_2, B_3 \subseteq \Theta_s, \quad M_3(B_3) = (1 - K)^{-1} \sum_{B_1 \cap B_2 = B_3} M_1(B_1)M_2(B_2),$$

$$\text{for } K = \sum_{B_1 \cap B_2 = \emptyset} M_1(B_1)M_2(B_2) \leq 1,$$

where K is the total amount of conflict between M_1 and M_2 , and $(1 - K)^{-1}$ is a renormalization factor.

Using Dempster's rule to combine evidence accomplishes two functions. The first is to obtain a consensus about what actions each CKS believes is possibly the best to take. For instance, one CKS might provide evidence that pertains to the costs associated with invoking various KSs, whereas a second distinct CKS may provide evidence that pertains to the reliability of invoking various KSs. Given the evidence they each provide, the task is to determine what actions both CKSs agree are the most appropriate to take. If both bodies of evidence are completely consistent, there is at least one action that they both agree should be pursued, and it can be said they are expressing totally compatible opinions—i.e., $K = 0$. Conversely, if the evidence they provide is not completely consistent, their opinions are then not completely compatible and there is at least one action the CKSs disagree is appropriate—i.e., $0 < K \leq 1$. In general, to the degree that CKSs are certain, complete, and accurate with respect to their observations, their opinions about which actions are appropriate, (i.e., which propositions are most likely to be true), will be compatible. Dempster's rule determines simultaneously if there are any actions that CKSs agree are appropriate and provides a measure of compatibility among the bodies of evidence they provide.

The second function of Dempster's rule is to correct for minor errors. The assumption here is that there is only a negligible likelihood that distinct CKSs

³ Shafer [SHA76] and Lowrance [LOW82b] show how Dempster's rule can be applied recursively to combine multiple bodies of evidence.

might introduce the same type of error into their bodies of evidence simultaneously. Therefore, such errors can be overcome by a sufficient amount of redundant and generally correct evidence. If a subset of CKSs make gross errors, such bad information should be discarded, when detected.

AN EVIDENTIAL APPROACH

The application of our approach begins by invoking a subset of the available CKSs so as to obtain their opinions about some aspect of the system's control environment, see Figure 5. In our approach, each CKS conveys its observations in the form of a mass distribution that expresses the CKS's belief in propositions such as the following:

- The risk, to the patient, of delayed treatment is great.
- Our goal should be to determine the patient's reaction to various medication before prescribing any medicine.

The mass distributions provided by the CKSs are then combined by using Dempster's rule [DEM67, DEM68]. Next the result of applying Dempster's rule must be extrapolated from those propositions on which it bears directly to the remaining dependent propositions. This extrapolation process is performed by an inference engine that is similar to the one described in Lowrance's thesis [LOW82b]. The inference engine performs its task by adjusting the bounds on the evidential interval that is associated with every dependent proposition. The result of the extrapolation process is a partial ordering over the set of alternative actions, as reflected in the evidential intervals associated with each $\alpha \in \Theta_a$. Selection of the appropriate action requires that these evidential intervals be evaluated. Although a complete decision theory for performing such an evaluation is not yet available, it is possible to choose actions on the basis of several simple criteria. For example, the best action is obvious for those propositions that correspond to alternatives with nonoverlapping evidential intervals. For those propositions with overlapping evidential intervals, further evaluation is called for. There are many utility- vs. cost-based theories that can be used to select an action on the basis of beliefs that are constrained by an evidential




Figure 5.

A DIAGRAM OF AN EVIDENTIALLY BASED APPROACH TO REASONING ABOUT CONTROL.

interval. Although the details of how such theories may be employed are beyond the scope of this paper, one approach involves invoking the appropriate CKS to evaluate the expected benefit and cost of pursuing each alternative. The CKS would express its opinion as to what alternatives it believes would be of greatest benefit, taking the cost of pursuing each alternative into account, in the form of a mass distribution. More mass will be attributed to those actions whose implementation the CKS believes would result in the greatest benefit.

A METHOD

We can construct P , a set of propositions, such that there exists a unique $p \in P$ for the subsets of Θ_a that the CKSs need to express their observations through or that are of interest. Our method for reasoning about control begins by having each CKS $_k$ express its belief about the elements of P through the mapping

$$M_k : 2^{\Theta_a} \mapsto [0, 1], \quad \text{where } M_k(\emptyset) = 0, \quad \text{and}$$

$$\sum_{p \in P} M_k(p) = 1.$$

Let a mass distribution be represented as a "mass vector" that is defined to be

$$MV_k = \langle \dots (p, M_k(p)) \dots \rangle, \quad p \in P, M_k(p) \geq 0.$$

A CKS $_k$ attributes a portion of its belief to a proposition p by including in its mass vector the pair

$$(p, M_k(p)), \quad \text{where } M_k(p) > 0.$$

Dempster's rule [DEM67, DEM68] is then used to combine these mass vectors, resulting in a new mass vector that is denoted by

$$MV_{1 \oplus 2 \oplus \dots \oplus k} = MV_1 \oplus MV_2 \oplus \dots \oplus MV_k.$$

This new distribution is the input to an inference engine that performs its task as previously described.

AN IMPLEMENTATION

A computational model that may be used to partially implement our approach is a dependency graph model of evidential support (DGMES) [LOW82b]. Dependency graphs, a variant of an inference net, are a formal representation of dependency relations—logical dependencies between propositions. A dependency graph consists of a set of propositions (nodes), a set of evidential intervals (node values) that constrain the system's belief in the propositions, and a set of dependency relations (propositional connectives and associated mapping functions) that specify how the *Spt* and *Pls* values of each evidential interval should be adjusted during the extrapolation process.

To implement our approach, using the DGMES model, every proposition of interest is represented as a node in a dependency graph. These propositions correspond to a hypothesis either about the presence or absence of control feature values or about the appropriate action to choose. The dependencies between control feature values and actions are expressed through logical connectives, also represented as nodes in a dependency graph. For example, pursuing a particular action *implies* that certain control features have been observed. This dependency can be represented by the dependency graph in Figure 6.

An example of how mutually exclusive propositions, such as the elements of A , can be represented as a dependency graph is shown in Figure 7. Finally, we can represent a disjunction of mutually exclusive propositions, for instance $f_1^1 \vee f_1^2 \vee \dots \vee f_1^k$ for $2 \leq k \leq |\mathcal{F}_1|$, by the dependency graph in Figure 8.

The simple dependency graphs in Figures 6 through 8 are sufficient for understanding the examples we will present in this section. However, a detailed explanation of the remaining four propositional connectives that may be used – if and only if (IFF node), negation (NOT node), conjunction (AND node), and disjunction (OR node) – can be found in [LOW82b].




Figure 6.

A SIMPLE DEPENDENCY GRAPH OF AN IMPLICATION RELATION (IF NODE). THAT IS, THE ACTION a IMPLIES THAT THE PROPOSITION f_i IS TRUE.

Dependency Graph Implementation of the Medical Example

In this section we construct a dependency graph that corresponds to the frame of action Θ_a , for our medical example, and show how partial orderings may be induced over a set of alternative actions (i.e., every $a \in A$) as a result of plausible observations by CKSs. The dependency graph representation of Θ_a that is used throughout the following examples is shown in Figures 9-12. For readability, the dependency graph has been illustrated in four figures. Each figure represents a portion (i.e., a subgraph) of the complete dependency graph that is obtained by merging each of the dependency graphs in Figures 9-12. The propositions corresponding to each node (e.g., a_1, f_3^2, \dots , etc.) in the graph are those from the previous section. We shall introduce each example within the context of the following hypothetical




Figure 7.

A SIMPLE DEPENDENCY GRAPH OF AN EXCLUSIVE RELATION (X NODE). FOR EXAMPLE, EACH $a \in A$ IS MUTUALLY EXCLUSIVE.

scenario.

Consider a situation in which a doctor is on duty in the outpatient section of a hospital. He has a set of four CKSs at his disposal. He may invoke a subset of these to obtain some of the information that might be needed to decide upon the most appropriate clinical treatment of a patient's illness. The four CKSs (i.e., CKS1, CKS2, CKS3, CKS4) convey their observations through each of our four control feature spaces of interest – DIAGNOSTIC-TESTS, MONETARY GOAL/SUBGOAL, TEST RELIABILITY, and RISK TO HUMAN LIFE – respectively.

EXAMPLE 1

While on duty, the doctor is summoned to the emergency room to examine a




Figure 8.

A SIMPLE DEPENDENCY GRAPH ILLUSTRATING AN EXCLUSIVE-OR RELATION (XOR NODE). FOR EXAMPLE, THE DISJOINT PROPOSITIONS $f_1^1, f_1^2, \dots, f_1^k$, WHERE $2 \leq K \leq |\mathcal{F}_1|$.

patient who has been complaining of soreness in their throat and periodic hot and cold flashes. As good medical practice dictates, the patient's needs must first be assessed to help determine whether any treatment is required immediately or could be delayed without incurring undue risks. The task of obtaining the pertinent information for this assessment belongs to CKS4. On the basis of observations of the patient's facial expressions, the absence of records about their medical history, and answers to the doctor's preliminary questions, let us suppose that his CKS4 is reasonably sure that the risk of delayed treatment is relatively low or moderate. This opinion might be conveyed in the form of the following mass vector MV_4 .

$$MV_4 = ((f_4^2 \vee f_4^3 \ .6)(\Theta_4 \ .4)).$$

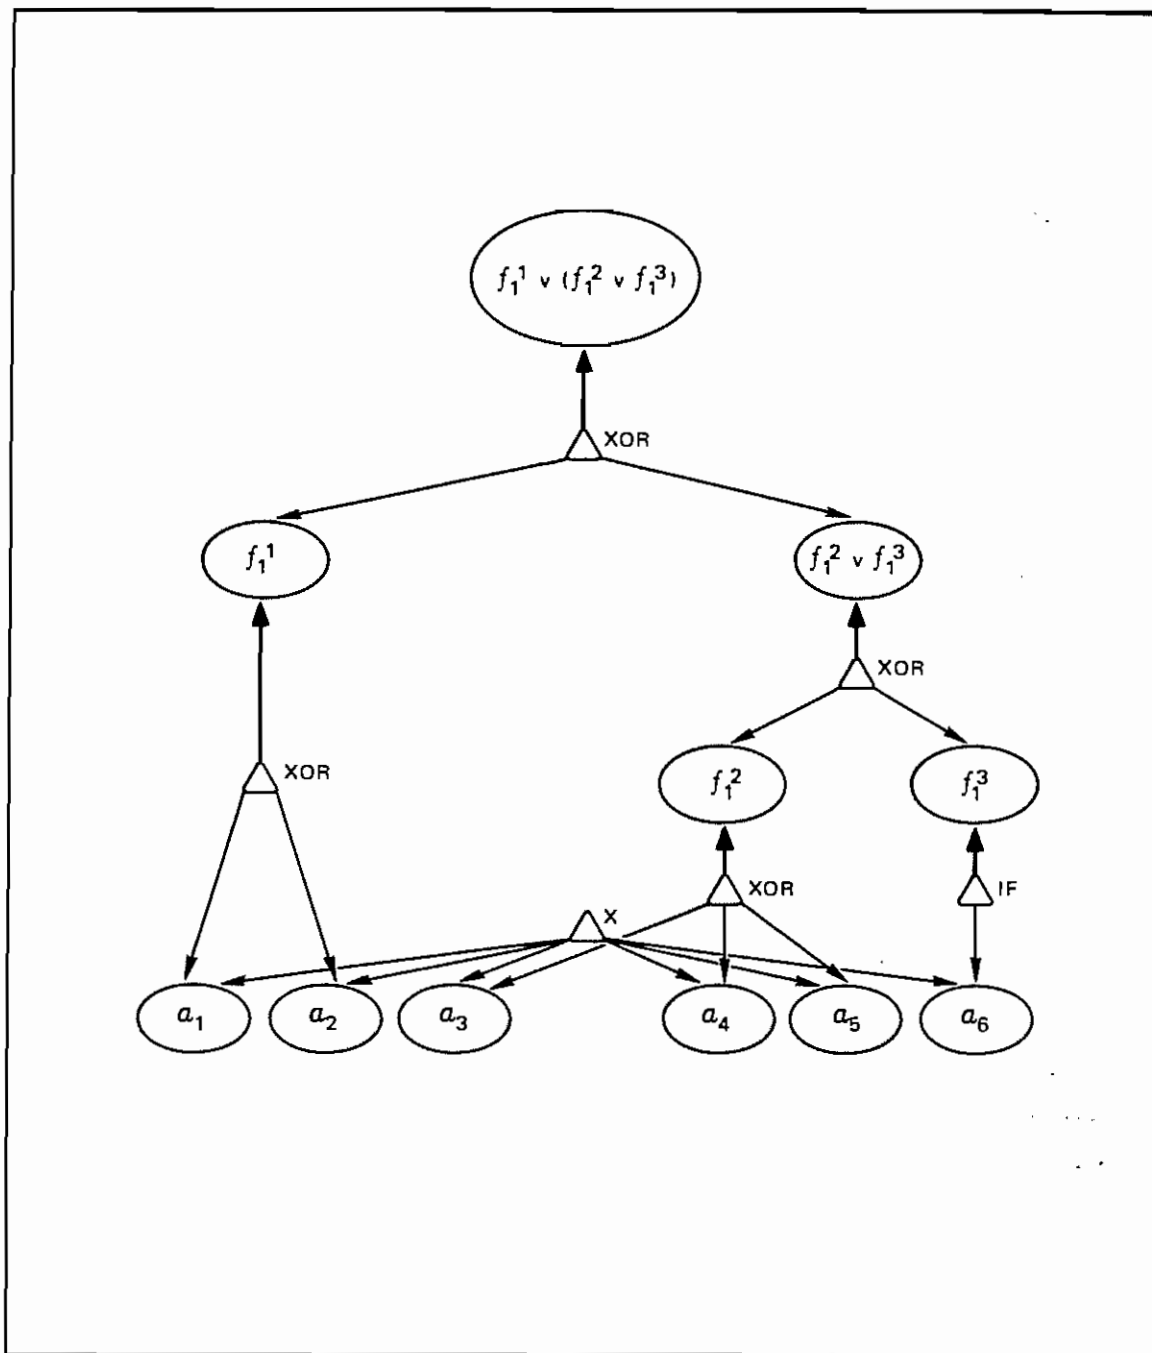


Figure 9.

A DEPENDENCY GRAPH REPRESENTATION OF THE DIAGNOSTIC TESTS CONTROL FEATURE SPACE F_1 AND ASSOCIATED VALUES f_1 .




Figure 10.

A DEPENDENCY GRAPH REPRESENTATION OF THE MONETARY GOAL/SUBGOAL CONTROL FEATURE SPACE F_2 AND ASSOCIATED VALUES f_2 .

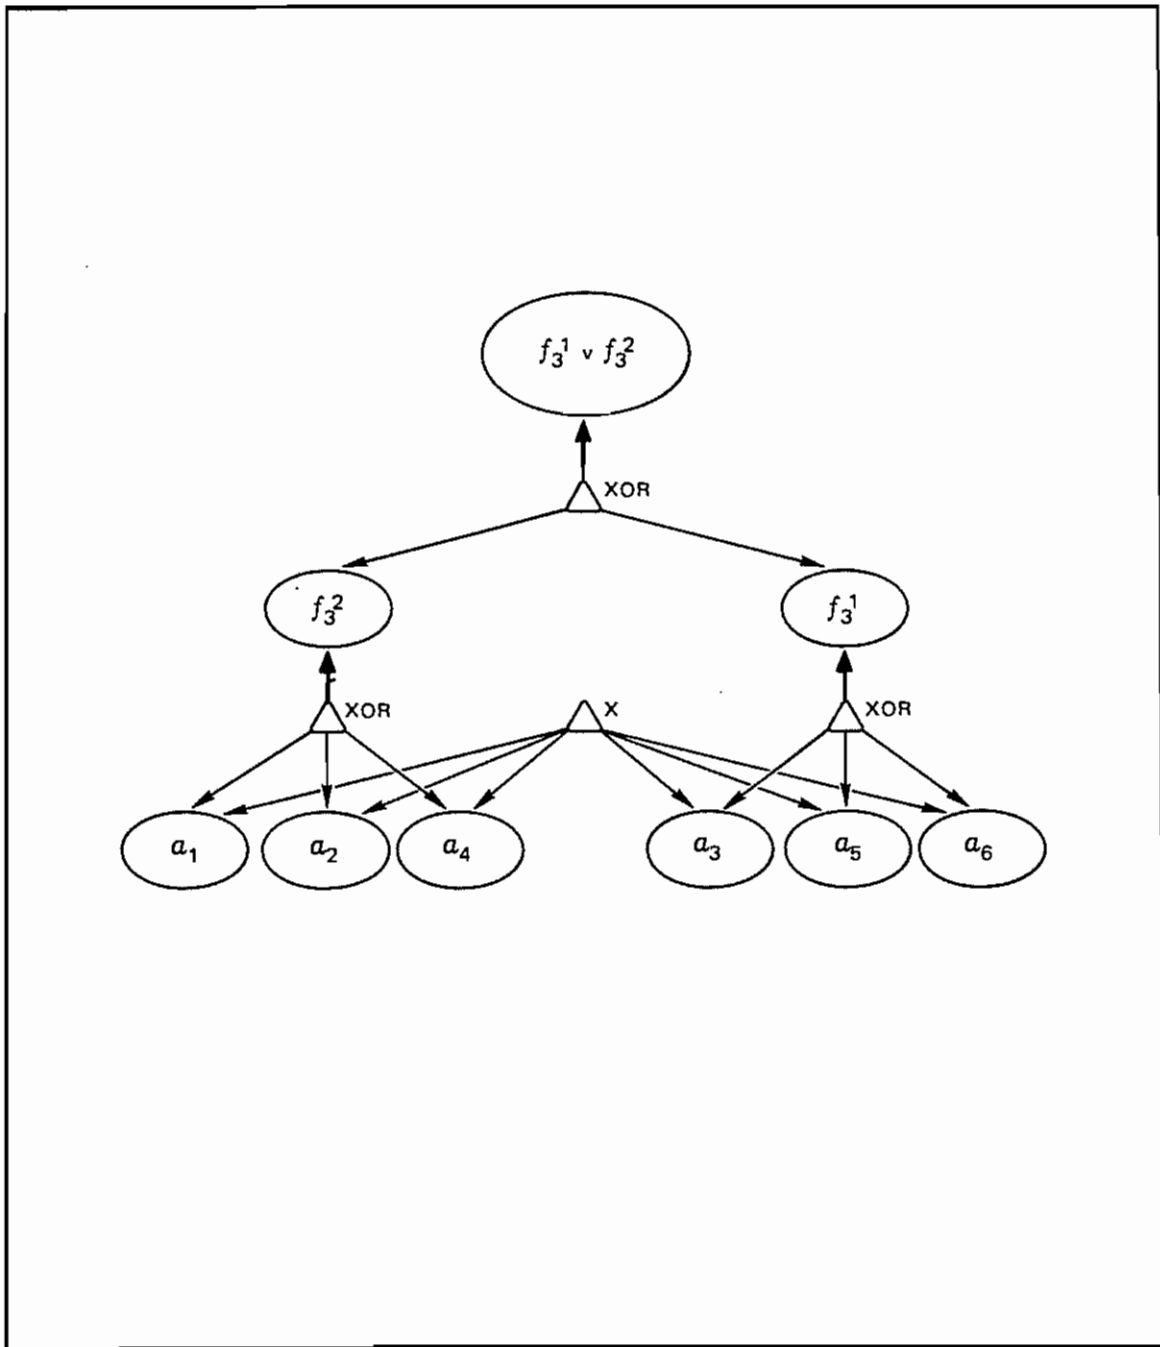


Figure 11.

A DEPENDENCY GRAPH REPRESENTATION OF THE TEST RELIABILITY CONTROL FEATURE SPACE F_3 AND ASSOCIATED VALUES f_3 .




Figure 12.

A DEPENDENCY GRAPH REPRESENTATION OF THE RISK TO HUMAN LIFE CONTROL FEATURE SPACE F_4 AND ASSOCIATED VALUES \mathcal{F}_4 .

Because the doctor is not absolutely sure that the risk of delayed treatment is low or moderate he is unwilling to commit all of his unit (i.e., 1) belief to any proposition. Therefore, a portion of his belief was attributed to Θ_a , which is logically equivalent to the disjunction of f_4^1 , f_4^2 , and f_4^3 . The result of extrapolating the effect of this opinion to our set of alternative actions is shown in Table 1. On the basis of CKS4's evidence alone, there is ambiguity about choosing any action. Actions a_3 , a_4 , a_5 , and a_6 , however, remain completely plausible. This indicates that, although there is no support for pursuing any alternative, some seem relatively implausible.

EXAMPLE 2

Ideally, most doctors are trained to treat a patient first—and only then to address the problem of how the financial obligation will be met. That is, the physician's goal/subgoal should be to make a diagnosis regardless of the monetary expense to the patient. However, reality may exert pressures that compel one to adopt opinions and exhibit behavior that are less than optimally desirable. Currently doctors are being strongly urged to avoid excessive diagnostic procedures and tests in their medical practice. Consequently, opinions about what monetary goals/subgoals ought to be satisfied, in conjunction with other relevant opinions, might be incorporated into the decision-making process. In this example, CKS2 has the task of obtaining the information necessary for forming such opinions.

Suppose the patient informs both the doctor and the hospital that he has limited medical coverage. While neither the doctor nor the hospital wish to deny treatment to a person in genuine need of medical attention, nor do they want to incur expenses beyond whatever may be absolutely necessary. Taking both the patient's medical coverage information into account and the current climate dictating prudent medical behavior, CKS2 might be relatively sure that the physician's goal/subgoal should be to make a diagnosis at moderate cost to the patient. This opinion might be expressed through the following mass vector MV_2 :

$$MV_2 = ((f_2^2 \quad .4)(\Theta_a \quad .6)).$$

This mass vector reflects the opinion that the doctor is relatively uncertain about

what his monetary goal/subgoal should be. Thus, most of his unit belief was attributed to Θ_a .

To obtain a consensus as to which alternative both CKS2 and CKS4 believe is the most appropriate one to pursue, we apply Dempster's rule to combine their mass vectors into a new mass vector $MV_{2\oplus 4}$:

$$MV_{2\oplus 4} = ((a_3 .08)(a_4 .08)(a_5 .08)(f_2^2 .16)(f_4^2 \vee f_4^3 .36)(\Theta_a .24)).$$

The effect of extrapolating the effect of the combined opinions to our alternatives is shown in Table 2. Although ambiguity remains about what action to take, CKS2's opinion, combined with that of CKS4, results in increased support for actions a_3 , a_4 , a_5 . In addition, pursuing actions a_1 , a_2 , and a_6 has become less plausible.

EXAMPLE 3

Opinions about the reliability of any test for the presence of bacteria or viruses can be influenced by many factors. Knowledge about the type of equipment on which the test is performed, the qualifications of the laboratory technicians who will conduct the tests, and even the physician's own personal bias about the level of quality control in the lab help to form opinions about the reliability of results.

After taking all these factors into account, let us assume that CKS3 is quite certain that the tests for bacteria are reliable. This opinion could be conveyed through the mass vector MV_3 :

$$MV_3 = ((f_3^1 .7)(\Theta_a .3)).$$

To obtain a consensus as to which alternative CKS2, CKS3, and CKS4 believe is the most appropriate, we apply Dempster's rule to the mass vectors $MV_{2\oplus 4}$

and MV_3 to produce a new mass vector $MV_{2\oplus3\oplus4}$:

$$MV_{2\oplus3\oplus4} = ((a_3 \ .248) \\ (a_4 \ .024) \\ (a_5 \ .248) \\ (a_6 \ .084) \\ (f_2^2 \ .048) \\ (f_3^1 \ .168) \\ (f_4^2 \vee f_4^3 \ .108) \\ (\Theta_a \ .072)).$$

The effect of extrapolating the new mass vector $MV_{2\oplus3\oplus4}$ is shown in Table 3. The results listed there indicate that some choices are clearly more appropriate than others. For example, since the evidential intervals for a_3 and a_5 do not overlap with the evidential intervals for a_1 and a_2 , if the physician had to choose between these two subsets, the choice should be $\{a_3, a_5\}$. However, there remains ambiguity among a_3 , a_4 , a_5 , and a_6 , which the doctor will have to resolve.

EXAMPLE 4

Even though monetary considerations might dictate treating a patient with no more than a minimum number of tests and diagnostic procedures, occasionally conflicting factors must be taken into account when a doctor is deciding how to diagnose his patient's illness. For instance, the patient might already be on medication. Uncertainty as to the exact cause of the illness might result in prescribing medicine that is incompatible with what the patient is currently taking. Additional tests, in the judgment of the doctor, may be necessary despite the fact that such tests have typically not been requested in analogous situations.

After reviewing the patient's medical records, let us suppose that CKS1 is reasonably certain that the physician should diagnose the patient's illness by requesting extensive tests. This opinion might be expressed through the mass vector MV_1 :

$$MV_1 = ((f_1^3 \ .6)(\Theta_a \ .4)).$$

We can obtain a consensus as to which action CKS1, CKS2, CKS3, and CKS4 believe is the most appropriate alternative to pursue by applying Dempster's rule to MV_1 and $MV_{2\oplus3\oplus4}$. This will result in a new mass vector $MV_{1\oplus2\oplus3\oplus4}$:

$$\begin{aligned}
 MV_{1\oplus2\oplus3\oplus4} = & ((a_3 \quad .13) \\
 & (a_4 \quad .013) \\
 & (a_5 \quad .13) \\
 & (a_6 \quad .461) \\
 & (f_1^3 \quad .057) \\
 & (f_2^2 \quad .025) \\
 & (f_3^1 \quad .088) \\
 & (f_4^2 \vee f_4^3 \quad .057) \\
 & (\Theta_a \quad .038)).
 \end{aligned}$$

The effect of CKS1's beliefs on our set of alternatives is shown in Table 4. According to the results so far, it is clear the doctor should take the following action: take a culture and request extensive tests for bacteria and viruses, then prescribe penicillin if bacteria tests are positive; otherwise if viral tests are positive, prescribe rest and take liquids.

COMPUTATIONAL ASPECTS

A large body of knowledge about control features and possible actions (commonly called metaknowledge) could be needed by a system. In some task domains, the representation of every control feature value as a node in a dependency graph might result in an immense number of nodes to be processed. Two methods for dealing with this predicament are to partition the dependency graph into conceptually meaningful and manageable segments or to consolidate a subset of the metaknowledge into fewer nodes.

An implied assumption of partitioning metaknowledge is that it will not be necessary for the system to observe the effect of evidence on all its knowledge at any one instant during the execution of its task. Therefore, only those partitions of the metaknowledge over which the system desires to draw inferences need to be

managed at any one time. A consequence of this technique is that the system must consider the possible effects of using only a subset of the potentially relevant control knowledge. The mechanisms needed to implement this method is a separate and important topic that is beyond the scope of this paper.

There are two conditions under which control knowledge may be consolidated. The first exists when a CKS can convey its observations adequately through a single proposition that corresponds to a disjunction of two or more control features. The second is when the system does not have to observe the effect of evidence on a subset of the control features that are represented by a single node. The consolidation of knowledge about control feature values in this manner reduces the computational overhead of drawing inferences. However, this benefit occurs at the expense of not being able to provide some highly desirable information. For example, the system will not be able to explain its choice of some actions on the basis of individual control features that have been represented as a single node.

The process of combining mass distributions requires that the conjunction of propositions be determined. The logical load of resolving various conjunctions depends on the complexity and completeness of the model. A possible solution whenever a conjunction cannot be resolved immediately has been proposed in [LOW82a]. That is, judgement can be suspended by reserving the appropriate portion of mass for that unresolved conjunct, thereby preventing it from contributing to either the support or plausibility of any proposition. If the conjunction should later be resolved, this mass can be redistributed in the appropriate places and allowed to influence the support and plausibility of other propositions.

Dempster's rule is both commutative and associative. This permits results to be obtained through hierarchical combination of partial results, with whatever degree of parallelism the host hardware can support.

CKS NAME	MASS VECTOR
CKS3	$MV_4 = ((f_1^2 \ .6)(\Theta_a \ .4))$
*** INFERENCE RESULTS ***	
proposition (node name)	evidential interval
a_1	[0, .4]
a_2	[0, .4]
a_3	[0, 1]
a_4	[0, 1]
a_5	[0, 1]
a_6	[0, 1]
f_1^1	[0, 1]
$f_1^1 \vee (f_1^2 \vee f_1^3) = \Theta_a$	[1, 1]
f_1^2	[0, 1]
$f_1^2 \vee f_1^3$	[0, 1]
f_1^3	[0, 1]
f_2^1	[0, 1]
$f_2^1 \vee f_2^2 \vee f_2^3 = \Theta_a$	[1, 1]
f_2^2	[0, 1]
f_2^3	[0, 1]
f_3^1	[0, 1]
$f_3^1 \vee f_3^2$	[0, 1]
f_3^2	[0, 1]
f_4^1	[0, .4]
$f_4^1 \vee (f_4^2 \vee f_4^3) = \Theta_a$	[1, 1]
f_4^2	[0, 1]
$f_4^2 \vee f_4^3$	[.6, 1]
f_4^3	[0, 1]

TABLE 1.

THE RESULTS OF EXTRAPOLATING CKS4'S MASS VECTOR, MV_4 .

CKS NAME	MASS VECTOR
CKS2 and CKS4	$MV_{2\oplus 4}$
*** INFERENCE RESULTS ***	
proposition (node name)	evidential interval
a_1	[0 , .24]
a_2	[0 , .24]
a_3	[.08 , .84]
a_4	[.08 , .84]
a_5	[.08 , .84]
a_6	[0 , .6]
f_1^1	[0 , .76]
$f_1^1 \vee (f_1^2 \vee f_1^3) = \Theta_a$	[1 , 1]
f_1^2	[.24 , 1]
$f_1^2 \vee f_1^3$	[.24 , 1]
f_1^3	[0 , .76]
f_2^1	[0 , .6]
$f_2^1 \vee f_2^2 \vee f_2^3 = \Theta_a$	[1 , 1]
f_2^2	[.4 , 1]
f_2^3	[0 , .6]
f_3^1	[.16 , .92]
$f_3^1 \vee f_3^2$	[.24 , 1]
f_3^2	[.08 , .84]
f_4^1	[0 , .4]
$f_4^1 \vee (f_4^2 \vee f_4^3) = \Theta_a$	[1 , 1]
f_4^2	[.16 , .92]
$f_4^2 \vee f_4^3$	[.6 , 1]
f_4^3	[.08 , .84]

TABLE 2.

THE RESULTS OF EXTRAPOLATING $MV_{2\oplus 4}$.

CKS NAME	MASS VECTOR
CKS2, CKS3, and CKS4	$MV_{2\oplus3\oplus4}$
*** INFERENCE RESULTS ***	
proposition (node name)	evidential interval
a_1	[0 , .072]
a_2	[0 , .072]
a_3	-.248 , .644]
a_4].024 , .252]
a_5	-.248 , .644]
a_6	-.084 , .432]
f_1^1	[0 , .396]
$f_1^1 \vee (f_1^2 \vee f_1^3) = \Theta_e$	[1 , 1]
f_1^2].52 , .916]
$f_1^2 \vee f_1^3$].604 , 1]
f_1^3].084 , .48]
f_2^1	[0 , .348]
$f_2^1 \vee f_2^2 \vee f_2^3 = \Theta_e$	[1 , 1]
f_2^2].568 , .916]
f_2^3].084 , .432]
f_3^1].748 , .976]
$f_3^1 \vee f_3^2$].772 , 1]
f_3^2].024 , .252]
f_4^1	[0 , .288]
$f_4^1 \vee (f_4^2 \vee f_4^3) = \Theta_e$	[1 , 1]
f_4^2].272 , .668]
$f_4^2 \vee f_4^3$].712 , 1]
f_4^3].332 , .728]

TABLE 3.

THE RESULTS OF EXTRAPOLATING $MV_{2\oplus3\oplus4}$.

CKS NAME	MASS VECTOR
CKS1, CKS2, CKS3, CKS4	$MV_{1\oplus 2\oplus 3\oplus 4}$
*** INFERENCE RESULTS ***	
proposition (node name)	evidential interval
a_1	[0 , .038]
a_2	[0 , .038]
a_3	[.131 , .339]
a_4	[.013 , .133]
a_5	[.131 , .339]
a_6	[.461 , .701]
f_1^1	[0 , .208]
$f_1^1 \vee (f_1^2 \vee f_1^3) = \Theta_4$	[1 , 1]
f_1^2	[.274 , .482]
$f_1^2 \vee f_1^3$	[.792 , 1]
f_1^3	[.518 , .726]
f_2^1	[0 , .24]
$f_2^1 \vee f_2^2 \vee f_2^3 = \Theta_4$	[1 , 1]
f_2^2	[.299 , .539]
f_2^2	[.461 , .701]
f_2^3	[.811 , .987]
$f_2^3 \vee f_2^4$	[.823 , 1]
f_2^4	[.013 , .189]
f_4^1	[0 , .208]
$f_4^1 \vee (f_4^2 \vee f_4^3) = \Theta_4$	[1 , 1]
f_4^2	[.143 , .408]
$f_4^2 \vee f_4^3$	[.792 , 1]
f_4^3	[.592 , .857]

TABLE 4.

THE RESULT OF EXTRAPOLATING THE MASS VECTOR
 $MV_{1\oplus 2\oplus 3\oplus 4}$.

FUTURE WORK

Future work will explore such an approach to control in several domains: general-purpose computer vision (e.g., the VISIONS system [HAN78]), and multisensor integration [GAR81]. Simultaneously, work will be done to develop an alternative representation of dependency graphs as a method for further reducing the computational load of resolving conjunctives and extrapolating mass throughout a dependency graph.

SUMMARY

Our interest in an evidential approach to control is motivated by its ability to address some important problems confronting expert systems that operate in complex domains. For example, some problems the underlying theory allows us to address are: (1)combine distinct types of evidential information from disparate sources; (2)provide a uniform representation of ignorance about the propositions of interest; (3)correct for minor errors in bodies of evidence; (4)could be used to reason about allocating additional resources in order to obtain the evidential information required to choose among alternatives; (5)facilitate a system's ability to explain its actions and; (6)flexibly order its actions by allowing any number of control strategies (i.e., control feature spaces) to be imposed over its alternatives. This approach is also modular with respect to the process of combining distinct bodies of evidence and extrapolating the result over a representation of metaknowledge. Finally, such an approach to control is conservative, because it does not suggest taking any action beyond that indicated by the evidence, a highly desirable attribute.

Although the concept of evidential reasoning and this related work is new, preliminary results confirm that this approach is promising with respect to its ability to deal adequately with some of the control-related issues and problems confronting expert systems.

ACKNOWLEDGEMENTS

We would like to thank Professors Allen R. Hanson, Edward M. Riseman, Charles Randall, Dr. Arthur Farley, and Thomas Strat for their extremely helpful suggestions.

REFERENCES

- [BAR81] Barnett, Jeffrey A., "Computational Methods for a Mathematical Theory of Evidence," Proc. Seventh International Joint Conference on Artificial Intelligence (IJCAI), University of British Columbia, Vancouver, B.C., Canada, pp 868-875, (24-28 August 1981).
- [BAR82] Barnett, J.A., "Some Issues of Control in Expert Systems," Proc. International Conference on Cybernetics and Society, pp 1-5., (October 1982).
- [DEM67] Dempster, A. P., "Upper and lower probabilities induced by a multivalued mapping," *Annals of Mathematical Statistics*, Vol. 38, pp. 325-339, (1967).
- [DEM68] Dempster, A.P., "A generalization of Bayesian inference," *Journal of the Royal Statistical Society, Series B*, Vol. 30, pp. 205-247, (1968).
- [GAR81] Garvey, T. D., J.D. Lowrance, M. A. Fischler, "An Inference Technique for Integrating Knowledge From Disparate Sources," Proc. Seventh IJCAI pp 319-325, (24-28 August 1981).
- [LOW76] Lowrance, J. D., "GRASPER 1.0 Reference Manual," COINS Technical Report 78-20, University of Massachusetts, Amherst, Massachusetts. (1976).
- [LOW79] Lowrance, J.D. and D. Corkill, "The Design of GRASPER 1.0: A Programming Language Extension for Graph Processing," COINS Technical Report 79-6, University of Massachusetts, Amherst, Massachusetts (1979).
- [LOW82a] Lowrance, J.D. and T.D. Garvey, "Evidential Reasoning: A Developing Concept," Proc. International Conference on Cybernetics and Society, October, pp 6-9 (1982).
- [LOW82b] Lowrance, J.D., "Dependency-Graph Models of Evidential Support," Ph.D. dissertation, Department of Computer and Information Science, University of Massachusetts, Amherst, Massachusetts., (September 1982).
- [SHA76] Shafer, G., *A Mathematical Theory of Evidence*, (Princeton University Press, Princeton, New Jersey, 1976).

- [SHO76] Shortliffe, E. H., *Computer-Based Medical Consultations: MYCIN*, (Elsevier Scientific Publishing Co., New York 1976).
- [WES82] Wesley, L., "The Use of an Evidential-Based Model for Representing Knowledge and Reasoning about Images in the VISIONS System," Proc. Workshop on Computer Vision: Representation and Control, Rindge, New Hampshire, pp 14-25, (August 1982).