A KNOWLEDGE-BASED DSS FOR MANAGERIAL PROBLEM DIAGNOSIS

James F. Courtney, Jr., and David B. Paradice
Business Analysis and Research Department, College of Business Administration,
Texas A&M University, College Station, TX 77843-4217
Nassar H. Ata Mohammed
King Faud University of Petroleum, Dhahran, Saudi Arabia

ABSTRACT

A knowledge-based system supporting managerial problem diagnosis is described. The system provides the capability to monitor values of selected variables for problem situations. When problems are located, a list of problem symptoms is delivered to a problem processor for structuring and diagnosis. Problem structuring is based on a combination of concepts from expert systems and structural modeling. User assertions about cause-effect relationships between pairs of variables are maintained in a semantic network. Problem diagnosis uses the relationships in the semantic network to construct causation trees, the branches of which represent potential explanations of the problem symptoms. Mathematical models are constructed based on causation-tree branches, and values from the data base are used to test whether the model confirms the diagnosis. If so, the source of the problem has been located and it is then up to the user to resolve the problem. If the model fails to explain the problem, the model apparently is deficient and the user may perform "what if. . ." type scenarios in attempts to improve the model and search for problem causes. Realistic applications in the accounting and health care areas are discussed.

Subject Areas: Computer Applications, Decision Processes, Decision Support Systems, and Management Control.

INTRODUCTION

Problem diagnosis refers to hypothesizing about or searching for causal relationships among variables believed to be associated with the problem at hand. The work described in this paper was motivated by the belief that problem diagnosis is a critical aspect of the managerial decision process, but one that has been neglected by decision support system (DSS) researchers.

Support for the belief that problem diagnosis is critical but neglected is given by several authors [83] [79] [6] [95]. As Mintzberg, Raisinghani, and Theoret put it,

> Diagnosis is probably the single most important routine [i.e., subprocess] since it determines in large part, however implicitly, the subsequent course of action. Yet researchers have paid almost no attention to diagnosis, preferring instead to concentrate on the selection routines. [83, p. 274]

In two independent surveys, DSS users rated problem diagnosis as one of the most important but least supported decision-making steps [79] [6]. This has led to the conclusion that existing techniques provide adequate support for problem analysis but very limited support for problem diagnosis [95].

The objective of the present authors' work has been to develop intelligent decision support tools for managerial problem diagnosis. In accordance with the perspective adopted in most DSS reseach, we have approached this problem from a strategic level. A major premise of our work is that to diagnose organizational problems from this perspective it is necessary to have a global representation of the organization's operating environment, both internal and external. We refer to this representation as the structure of the organization's problem domain or the global problem domain. It is important to represent the structure of the organization's problem domain because problems may be caused by any element within that domain—either inside or outside the organization itself. Thus, if problem diagnosis is to be supported at the strategic level, knowledge of the global problem domain must be obtained, represented, and managed by the system.

The next section of our paper describes the problem of managerial problem diagnosis more fully. Artificial intelligence (AI) research on problem diagnosis, reviewed in the third section, indicates that diagnostic systems are more effective when based on causal ("deep") structure. Research on causal structure in management problems is found in the literature on model management systems (MMS) and structural modeling, which also is reviewed to show its relevance to managerial problem diagnosis. Then concepts from these different areas are integrated and an approach that provides computer-based support for managerial problem diagnosis is developed. Finally, some applications are discussed and conclusions are drawn.

MANAGERIAL PROBLEM DIAGNOSIS

A number of models describing decision processes have been presented in the literature (see, e.g., [106], [83], [58], [3], [4], [112], and [102]). Work in DSS is concerned with providing computer-based assistance for one or more phases or steps of such managerial decision processes, especially in ill-structured decision situations. Since ill-structured situations are common at the strategic level of an organization, much DSS research has been concerned with support for strategic planning and decision making at the organizational apex.

Most models of the managerial decision process begin with some type of environmental scanning or problem finding activity. Pounds [93] observed that managers "find" problems when observed conditions do not correspond to anticipated or expected conditions.

Once a problem has been located, the next step typically is described as problem structuring [95], problem formulation [5], or problem diagnosis [83]. The terms "structuring," "formulation," and "diagnosis" are used to refer to the same or very similar concepts. For example, problem diagnosis has been defined both as hypothesizing about cause-effect relationships between variables involved in the problem [83] and as "backward inference" from observed symptoms and signs to prior causes [41]. Problem structure has been defined as involving the elements or variables in a problem and how those variables fit together and interact [1]. Problem formulation is similar but may be restricted to a mathematical representation of the problem [102] [5].

Intuitively, problem diagnosis, structuring, and formulation are closely related because each involves the specification of causal relationships among variables. It seems unlikely that a problem can be diagnosed without at least a hypothesis about its structure. Indeed, hypothesizing about cause-effect relationships *is* hypothesizing about structure [78]. Furthermore, mathematical representations of a problem are based on some understanding of the problem's structure. Thus, even though the work in this paper is directed toward problem diagnosis, it is related directly to problem structuring and formulation since improper formulation may arise from incorrect diagnoses [97] [84] [120].

Einhorn and Hogarth [41] discussed causal relationships and diagnosis, emphasizing the fact that diagnosis and prediction are based on the *same* causal structure. Thus, causal structure is useful both in diagnosis and in planning since planning inevitably involves predicting the outcome of possible courses of action.

Einhorn and Hogarth used a 2×2 table as a framework for analyzing how people infer causal relationships between two variables, X and Y. The rows of their table are labeled X and $\bar{X}$ (not-X, indicating the assumption that X did not occur) and the columns are Y and $\bar{Y}$. Six different types of cues are used to determine whether or not a causal relationship exists: temporal order, contiguity in space and time, constant conjunction, number of alternative explanations, similarity, and predictive ability.

Temporal order refers to which of the two variables occurred first, X or Y. This is used to determine which is the potential cause and which the effect. In the ensuing discussion we will assume X occurred before Y, so we are attempting to determine whether X may have caused Y.

Contiguity in space and time refer to the notion that people have a stronger tendency to feel X caused Y if X and Y are close together in space and time. If so, then the strength or "frequency" (as Einhorn and Hogarth called it) of the cell corresponding to X causing Y is greater relative to the other cells.

Constant conjunction is "the degree to which X and Y occur together, holding contiguity in space and time constant" [41, p. 27]. This also relates to the relative strength of cell X,Y.

The number of alternative explanations refers to cell $\bar{X},Y$ as it asks the question, "Would Y have occurred if X had not?" The greater the number of ways that Y may have occurred in the absence of X, the lower the causal relevance of X to Y and the lower the relative strength of cell $\bar{X},Y$.

Similarity occurs within the context of some metaphorical analogy, such as comparing the human brain to a computer. It is defined as a weighted linear function of confirming evidence (cells, X,Y and $\bar{X},\bar{Y}$) and disconfirming evidence (cells $\bar{X},Y$ and $X,\bar{Y}$).

The final cue, predictive ability, is the correlation coefficient between X and Y, defined over all four cells. The meaning of the correlation coefficient is *interpreted* in light of information about *other* cues, however. If other cues do not indicate sufficient reason to believe a causal relationship exists, the correlation is spurious.

We refer to these six causal cues when we discuss the experiential knowledge base (EKB) of the diagnostic system we propose and show how information about cues can be generated. In the next section, we describe research in AI and MMS that relates to diagnosis.

RELATED RESEARCH

Research related to managerial problem diagnosis and problem structuring is found in the AI and MMS literatures.

AI Research on Diagnosis

Stefik, Aikins, Balzer, Benoit, Birnbaum, Hayes-Roth, and Sacerdoti [107] indicated that (nonmanagerial) expert systems have been developed for a wide variety of tasks including planning, design, prediction, diagnosis, monitoring, and interpretation. They referred to systems that perform more than one of these tasks as "control systems." All six tasks occur in the management domain. Our main concern in this paper is with managerial problem diagnosis; however, we later show how this task relates to other tasks as well.

The crux of diagnostic activity is the ability to infer system malfunctions from observed activities and to relate the observed behavioral irregularities to underlying causes [24]. Stefik et al. [107] defined diagnosis as a process of faultfinding in a system based on interpreting potentially "noisy" data. Diagnosis requires understanding the system organization and the relationships and interactions among subsystems. Key problems are (1) faults can be masked by the symptoms of other faults; (2) faults can be intermittent; (3) diagnostic equipment can fail; (4) some data about a system are inaccessible, expensive, or dangerous to retrieve.

There has been a great deal of research related to *nonmanagerial* problem diagnosis in the field of AI. Many expert systems dealing with diagnosis in nonmanagerial domains have been constructed including INTERNIST and CADUCEUS (medicine) [92], MYCIN (medicine) [105], Alven (medicine) [86], PROSPECTOR (mineral exploration) [49], DIPMETER ADVISOR (geologic signals) [35], and an unnamed system for diagnosing grain combine malfunctions [103]. Further, a host of other diagnostic systems generally are known to researchers (see, e.g., [47]) including DART from IBM (computer systems faults) and PUFF from Stanford (lung disorders).

The unique features of these systems include (1) separation of knowledge about the problem domain into a knowledge base that can be manipulated as a separate entity, (2) an inference engine that operates on the knowledge base to extract knowledge and make inferences in response to user problem statements, (3) the use of probabilistic or quasiprobabilistic factors (certainty factors) concerning relationships or rules in the knowledge base to provide a way of dealing with problems that involve uncertainty, and (4) explanation facilities that are used to describe the reasoning the system uses to arrive at particular conclusions.

Recently, expert systems development has been based on inference engines that reason from "first principles" [33] [34] (the fundamental laws or relationships of the problem domain). Diagnosis from first principles also has been explored in a number of domains including medicine (CASNET [119], ABEL [91], RX [17]), computer-aided instruction [25], electronic troubleshooting [36], and physics [37] [69] [48].

Notably, all these systems have been developed in "deep and narrow" problem domains [13]. A deep and narrow problem domain is one where experts exist and the underlying laws of the problem domain are well formulated and usually static. Management, on the other hand, is a "wide and shallow" domain [13]. Managerial expertise generally is drawn from several sources and the fundamental laws frequently are ill-structured and even changing [39] [88].

Gentner and Stevens [51] confirmed that deep and narrow domains have been the most fruitful areas for research aimed at capturing models of human expertise. For example, recognizing an expert in physics is much easier than recognizing an expert in human relations. Furthermore, an expert's "mental model" [51] in a domain such as physics can be compared to the well-formed laws of the domain; such a comparison for an expert's mental model in a wide and shallow field is very difficult, if not impossible.

Although the management domain is wide and shallow, managerial problem diagnosis shares a fundamental characteristic with approaches taken in other domains. Specifically, causal explanations frequently are used to diagnose problems in managerial domains [5] [15]. Kasper [62], Pracht [95], and Loy [75] demonstrated the effectiveness of supporting the construction of mental models in a managerial problem domain. Paradice and Courtney [88] [90] and Pracht and Courtney [96] showed preliminary ways to capture, organize, and manipulate managerial mental models. The work reported herein seeks to combine recent insights gained in the field of AI with prior efforts in managerial problem formulation to support diagnosis of managerial problems—a significantly different type of problem domain than those addressed in prior studies.

Davis [34, p. 404] argued that models of causal interaction in diagnostic systems are superior to empirically based diagnostic systems. Such models are called "deep models" [53] [81] [28] [29] because they attempt to capture relationships and first principles that are at the basis of expert knowledge. Deep models have the capability to deduce behavior from structure [70] and may be able to build predictive models of future behavior [119].

Chandrasekaran and Mittal [29] observed that the straightforward approach taken by expert systems employing condition-action production rules is not feasible in problem domains of any significant size. Hence, deep models will be necessary in large problem domains. Bouwman's [23] work on financial problem diagnosis and research in MMS provides a partial basis for a diagnostic system for the wide and shallow domain of management, as described next.

Bouwman's Study of Financial Diagnosis

Bouwman [23] is one of very few researchers who has studied managerial problem diagnosis. In what may be regarded as landmark research, he used protocol analysis to study how financial analysts diagnose accounting statements. He then developed a computer program to mimic this diagnostic process in a qualitative manner. Bouwman's analysis revealed problem detection as the first step of the diagnostic process. Five operators were used in the problem detection process: computation of a simple trend, computation of a more complex trend, comparison with other information, comparison with a norm, and application of a heuristic or "rule of thumb." He also established rules for determining which operator to apply in a particular situation. These operators were used to translate the problem into qualitative terms.

The next step was to screen out insignificant findings, which his subjects then excluded from further analysis. Then came actual diagnostic reasoning, which involved integrating new findings with existing knowledge, and forming hypotheses that attempted to explain these findings. Existing knowledge used by subjects involved a model representing "a causal structure that describes the internal model of a typical firm" [23, p. 659]. Bouwman found that causal chains were integrated into "trees of causes," the branches of which represent different problem hypotheses or potential explanations of the problem's causes. Subjects mentally explored the branches of these trees in order to diagnose the problem. Bouwman was able to construct a sophisticated computer program that simulated this complex process very effectively. He even was able to tune his program to simulate the processes used by individual subjects.

From our perspective, one of the most important aspects of Bouwman's program is its reliance on the "internal model of a typical firm" to perform the diagnosis. This model consists of causal relationships among variables in the problem domain and was derived from the protocol analysis. The internal model was used to develop trees of causes (or causation trees, as we shall call them) on which diagnoses were based. Since the internal model is used to drive the construction of causation trees and the diagnostic process, the accuracy of the diagnosis is determined largely by the accuracy of the internal model. Thus, it is critical that the internal model accurately reflect the actual firm being modeled. We illustrate how the causation trees are constructed when our diagnostic system is described.

As in most DSS research, the concern in this work is with the development of tools to assist in formulating and using quantitative models. Bouwman's qualitative system provides an excellent starting point for a quantitative approach. In the next section we describe related work on problem structuring in MMS.

Model Management Systems (MMS)

Work on MMS is relevant to problem diagnosis because it deals with quantitatively oriented problem structuring. MMS seek to support dynamic problem structuring and to facilitate the use of mathematical models in managerial decision making [42] [44] [43]. As Lenard noted, "model management begins with a scheme for representing models and must provide for generating, restructuring, updating, and obtaining results of models" [72, p. 36]. Thus, research in MMS typically has treated models as entities and employed some knowledge representation technique to represent either relationships between variables within models (intramodel relationships) or between the models themselves (intermodel relationships) or both [44] [66] [82] [46].

Another branch of MMS research involves the use of techniques from database systems, often combining these with a knowledge representation scheme [18] [19] [20] [21] [22] [15] [16] [104] [115] [116] [72]. Again, models are the entities on which these approaches are based.

Our concern is with representing the structure of problem domains. Since MMS deal with problem structuring, it appears the MMS approach may be relevant to problem diagnosis. This approach does have shortcomings, however, because of the somewhat limited perspective adopted in MMS work.

Although Elam, Henderson, and Miller [44] suggested that MMS be able to incorporate broad knowledge of the problem domain, virtually all the MMS work cited above is limited to that portion of the domain for which models already exist. Since managerial problems may result from elements that are not contained in existing models, it is necessary to extend the techniques of MMS to the broader concept originally proposed by Elam et al. [44]. The next section describes how concepts from AI, MMS, and structural modeling may be integrated to form a more comprehensive approach to managerial problem diagnosis.

TECHNIQUES FOR REPRESENTING SYSTEM STRUCTURE

Techniques for representing and analyzing the structure of problems have been developed in a branch of systems engineering known as structural modeling. The seminal theoretical work on structural modeling is Harary, Norman, and Cartwright's [52] and relies heavily on the theory of directed graphs (digraphs). In structural modeling, problems have both a graphical and numeric (matrix) representation. Kane [61] referred to techniques that deal mostly with structural properties as geometric; those that deal with numeric representations he called arithmetic. For reasons discussed shortly, we chose to use the structural modeling approach to represent the global structure of problem domains.

A rather simple structural model is shown in Figure 1 for purposes of illustration. This model is meant to represent the global problem domain. Causation trees for specific problem instances will be derived from the global structural model. Causation trees have a particular problem variable as the root node, and directly related and indirectly related variables as children. Hence a causation tree is just that portion of a structural model relevant to a given problem.

Information in a structural model such as this is obtained from a person (or persons) knowledgeable in the problem domain—hopefully an "expert" in that domain. The information normally is collected via assertions about the relationship between pairs of variables in the domain. These assertions may take various forms such as

Changes in x_i cause changes in x_j,

Changes in x_i cause increases (or decreases) in x_j, or

A one-unit change in x_i causes x_j to change by c_{ij}.

The information in Figure 1 may be represented in a matrix by forming a row and column for each variable in the problem. The resulting matrix is called a binary connection matrix or adjacency matrix because a 1 is entered in cell i,j if there is a causal relationship from variable x_i to x_j; otherwise a zero is entered. The first form of assertion given above yields a binary adjacency matrix.

Figure 1: Structural model relating net earnings to marketing and production variables.

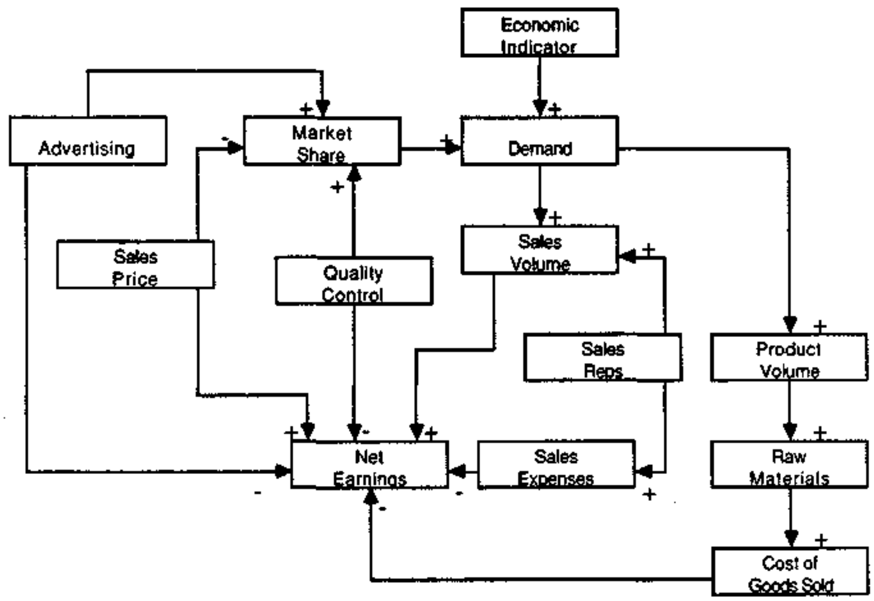

Matrix powering may be applied to adjacency matrices to determine paths of various lengths between variables. Specifically, if **A** is the adjacency matrix, then $\mathbf{A}^n$ gives all paths of length n in the matrix. Each element of $\mathbf{A}^n$ gives the total number of paths from variable i to variable j. A "total connection matrix" **T** gives the total number of paths between each variable pair and is given by: $\mathbf{T}=\Sigma\mathbf{A}^n$ for $n=1, 2, \ldots, p$, where p is maximum power to which **A** is raised.

From the point of view of problem diagnosis, matrix powering and the total connection matrix may be used to determine the paths from problem variables (problem symptoms) to potential problem causes. McLean and Shepherd [77] referred to this as qualitative information in the sense that it gives the *number* of paths between a problem and a cause, but not the *strength* of the cause.

Signs indicating the direction of change in x_j induced by the changes in x_i (the second form of assertion given above) may be added to the adjacency matrix to yield a signed digraph. If changes in x_i lead to increases in x_j then x_i is said to augment [100] or excite [26] [27] x_j. If x_i decreases x_j the relationship is said to be inhibitory. Burns and Winstead [26] [27] developed a geometric approach to determine the redundancy of paths in a model based on the excitatory or inhibitory nature of the path. Roberts [99] used signed digraphs to study energy demand.

In managerial problem diagnosis, some paths from causes to symptoms may be excitatory, others inhibitory. Thus the question we are concerned with is the net impact of the excitatory and inhibitory paths. In order to answer this question, information about the strength of the path must be given. Hence McLean and Sheperd's [77] observation that the strength of relationships often is critical is relevant for quantitative managerial problem diagnosis. Paradice and Courtney [90] showed how this information may be used in a managerial context.

If coefficients indicating the amount and direction of change induced in x_i by a one-unit change in x_j are included in the matrix (the third assertion form), then a "weighted adjacency matrix" is produced. Suppose, for example, we have a three-variable problem of the form $x_1 = f(x_2)$; $x_2 = f(x_1, x_3)$; $x_3 = f(x_1)$. Letting

$$C_{ij} = \frac{\partial x_i}{\partial x_j}$$

and writing these in matrix form yields an "interaction matrix" [77]. If each function is linear, each C_{ij} is a scalar; thus a weighted adjacency matrix is produced. Therefore the weighted adjacency matrix represents first-order changes in the relationships among variables in the model.

The arithmetic property of weighted adjacency matrices that we are concerned with is referred to as a pulse process [100] or pulse analysis [77]. This is a form of sensitivity analysis in which input variables are changed and the resulting change (or "pulse") in an output variable is computed over a given period of time.

Roberts [100] assumed discrete time periods and showed that the one-period pulse in variable i at time $(t+1)$ is given by

$$P_i(t+1) = \sum_j C_{ji} P_j(t).$$

He went on to show that *forecasts* of future values may be obtained from

$$\mathbf{V(t)} = \mathbf{V(0)} + \mathbf{P(0)}(\mathbf{A} + \mathbf{A}^1 + \ldots + \mathbf{A}^t)$$

where $\mathbf{V(t)}$ is the forecasted matrix of values at time t, $\mathbf{V(0)}$ is the initial value matrix, $\mathbf{P(0)}$ is the matrix of initial pulses, and $\mathbf{A}$ is the weighted adjacency matrix. Thus the capability of forecasting is inherent in the structural modeling approach.

This information on structural modeling provides the basis for our approach to managerial problem diagnosis. Other arithmetic techniques related to weighted adjacency matrices are described by Axelrod [11], Kruskal [67], McLean and Shepherd [77], McLean [76], Roberts [100], Waller [114], and Lendaris [73]. As a final remark, we note that several software packages have been developed for various structural modeling techniques. Some of the more popular are KSIM [61] [68], QSIM [113], SPIN [78], ISM [117] [118] [9], SMGS [71], GRIPS [56], and GISMO [94] [95].

The next section of the paper develops an approach to managerial problem diagnosis that integrates the pulse process technique of structural modeling with Bouwman's findings on managerial problem diagnosis. This approach extends Bouwman's approach in several ways.

PROPOSED APPROACH

Before describing our system for managerial problem diagnosis, it will be useful to describe the managerial decision-making context within which the system is presumed to operate. We describe a decision-making model that is not intended to be new or novel; it is used to set the context for subsequent discussions.

Our diagnostic approach is designed to support ongoing managerial control activity. The model we present (Figure 2) is similar to several in the literature [106] [83] [102] [58] but most closely approximates that of Ackoff [3]. Note that the model is cyclical in nature, emphasizing the ongoing nature of management control. The terminology of Stefik et al. [107] was chosen specifically to emphasize the relationship of the model to work in AI. Note that all the tasks in the model involve complex cognitive processes, many of which have been studied extensively in cognitive psychology. Some references to that literature are given where appropriate. We begin our discussion of the model with monitoring, the topmost rectangle in Figure 2.

Figure 2: Support for the management control process via a structural modeling knowledge base (SMKB) and an experiential knowledge base (EKB).

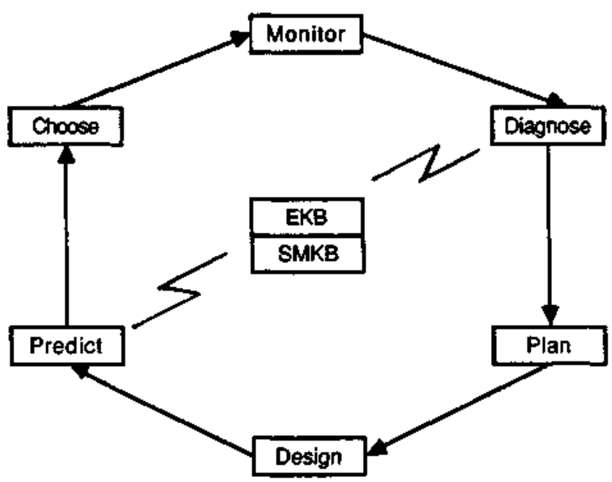

Monitoring involves acquiring information (perhaps via a data-base management system) describing the current state of the organizational system, comparing the current state to the desired (goal) state, and noting significant discrepancies [93] [41]. As Bouwman [23] found, these discrepancies between what is expected or desired and what is obtained constitute a list of significant findings or problem symptoms that is passed forward for diagnosis. Hogarth and Makridakis [55] listed six different types of human bias that may occur at this stage of the process. They suggested that a DSS should attempt to control the effect of such biases, some of which we describe briefly later.

Diagnosis is the next step and consists of searching for or hypothesizing about causes of the observed discrepancies or problem symptoms. Mintzberg et al. [83], Ackoff and Sasieni [5], and Bouwman [23] described managerial diagnostic activity. As discussed earlier, Einhorn and Hogarth [41] described how people decide whether or not a causal relationship exists between two variables. At least 10 different types of bias may affect information processing at this stage [55].

The next three steps, *planning, design,* and *prediction*, frequently form a feedback loop within the overall cycle [83]. Stefik et al. [107] described planning as setting new goals in light of the current situation. In our case, the problems confronting the organization are especially important, as are the resources available to attack these problems. Design consists of generating alternative ways of achieving the goals; prediction refers to forecasting the outcomes of each alternative.

Ackoff [4] described the tradeoffs between prediction and planning. That is, if we can predict well, then planning for contingencies is less important. Similarly, if our plans provide for many contingencies, then the need for accurate prediction is reduced. Einhorn and Hogarth [41] discussed the fact that causal models may be used in both diagnosis and forecasting. They noted that planning often leads to an unjustified "illusion of control" over an uncertain environment.

The final step in this cycle involves comparing the predicted outcomes of the various alternatives, *choosing* an alternative (presumably the one most likely to achieve the planned goals), and implementing it. Then new information is acquired and the cycle is repeated. Several types of bias also are associated with choosing an alternative, and several occur in a feedback process such as the one described [55].

We would like to emphasize that as the cycle is repeated, a learning process occurs within individual managers and the organization as a whole [63] [30] [96]. This learning process relates to gaining an understanding of the organizational problem domain [96]. Learning is another area of study in cognitive science that has been suggested as a "reference discipline" [64] for DSS [54] [65] [57] [102]. Our system includes a module that accumulates knowledge gained from the diagnostic process and thus is a form of experiential knowledge [65]. This is a unique aspect of our system. As illustrated later, it has uses in helping to control biases which may arise during the decision-making process. To emphasize the fact that structural knowledge of causal relations may be used in diagnosing, analyzing, and predicting problems, we have placed a structural modeling knowledge base (SMKB) in the center of Figure 2 with arrows indicating which processes it may support. We also have shown an experiential knowledge base (EKB) that also is useful in these processes.

ARCHITECTURE OF THE DIAGNOSTIC SYSTEM

The architecture of a system to support managerial problem diagnosis is illustrated in Figure 3. The full architecture supports both formal and concrete thought processes by including a module for storing the results ("experiences") of formal diagnostic analyses in the EKB, which serves as a memory for problems that have been diagnosed in the past.

Basic elements of the system include (1) a user interface, (2) a monitor that searches for problem situations, (3) a problem processor that searches for a problem's causes, (4) a knowledge manager that maintains causal relationships and historical information about successful diagnosis (the EKB), (5) a data manager that supplies data to the rest of the system as needed, (6) a dictionary manager that maintains information on all data items and variables in the system, and (7) a process-control subsystem that links various elements of the system together and controls the flow of information. This paper is concerned primarily with the problem processor and the knowledge bases, which are described in some detail.

The Monitor and List of Problem Symptoms

Since problems must be located before they can be diagnosed, the system includes a module (the monitor) for problem finding. Our approach to problem finding is based on the work of Pounds [93] who found that managers define problems as the difference between the present situation as observed and some expectation or standard. Expectations and standards may be based on recent trends, projections, or the behavior of comparable organizations. They may be the result of formal studies or derived from formal models; or they may be very informal and reside only in the manager's mind.

Since our primary concern is with problem diagnosis and not problem finding per se, our treatment of problem finding is very limited. The problem-finding module itself could be very complex. Sophisticated forecasting models (perhaps based on structural modeling techniques) could be used to predict anticipated values for selected variables (monitored variables) in the problem domain.

We assume the user selects monitored variables from among the variables in the problem domain and that standards somehow have been established for these monitored variables. The monitor's function simply is to compare observed values (the most recent values in the data base) to the standards for each monitored variable. Any variable violating its standards will be referred to as a problem symptom. The monitor prepares a list of problem symptoms and delivers this list to the problem processor for diagnosis.

Since organizational problem domains are highly complex and interrelated, it is possible that a single underlying cause may result in several problem symptoms. For example, a downturn in the economy may result in decreased sales and increases in inventory holding costs due to the declining demand. The problem processor must be intelligent enough to account for this phenomenon and to avoid excessive searches when one or a few causes have led to the discovery of several related problem symptoms.

The Problem Processor: Problem Structuring and Diagnosis

Support for problem diagnosis is based on structural modeling. As described previously, structural models may be based on assertions of the form: "A one-unit

Figure 3: Basic system architecture.

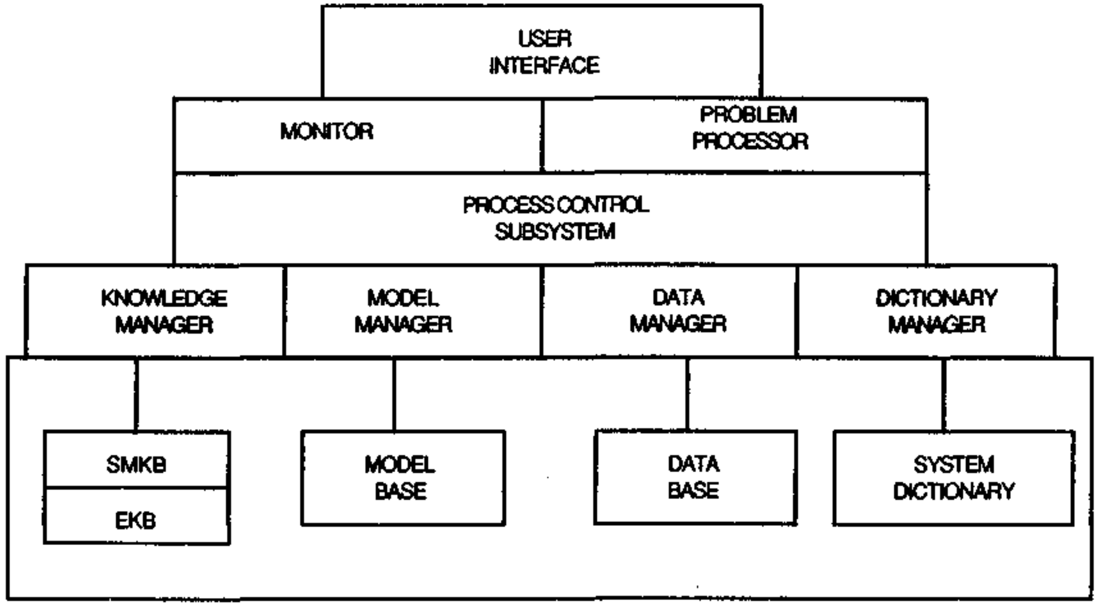

change in variable i causes variable j to change by c_{ij}." Such assertions are captured and stored in the knowledge base as a semantic network. Thus, as noted previously, the system begins with relationships between pairs of variables, not complete models as in MMS. As illustrated below, these relationships are combined to form models on which the diagnostic process is based. Relationships are based on changes observed in values between two points in time; thus this is a differential approach.

The following notation will be used to describe the diagnostic process:

X $=\{x_1, x_2, \ldots, x_n\}$, the indexed set of elements (variables) in the problem domain,

M = the set of monitored variables and is a subset of X,

S_t = the set of problem symptoms (monitored variables violating standards) at time t and is a subset of M,

$\Delta(x_i)$ = the observed change in x_i between times $t-1$ and t,

$\Delta'(x_i)$ = the change in x_i computed using the structural model currently stored in the knowledge base (how $\Delta'(x_i)$ is computed is described later).

Problem finding and diagnosis is a phasewise process involving (1) application of monitoring procedures to determine the set of problem symptoms, (2) use of Bouwman's approach and information in the knowledge base to construct clusters of related symptoms, (3) use of Bouwman's approach to construct a "causation tree" for each cluster, (4) construction of explanatory models based on the causation tree and extraction of data from the data base to test each model, and (5) presentation of results to the user.

Phase 1: Apply monitoring procedures to the elements of M, placing any variables violating their standards on S_t, the set of problem symptoms.

Phase 2: Apply Bouwman's approach to form clusters of related elements of S_t. Notice that clusters consist only of elements of S_t. Conceptually, related problem symptoms form one "problem" with the same underlying set of base causes. Thus a diagnosis will be attempted for each cluster. The following recursive definition of a cluster with symptom x_i as the root node may be implemented easily in a logic programming language such as PROLOG:

$$x_j \text{ clusters with } x_i \qquad \text{if } x_j \text{ impacts } x_i \text{ and } x_i \in S_t \text{ and } x_j \in S_t. \tag{1a}$$

$$x_j \text{ clusters with } x_i \qquad \text{if } x_k \text{ impacts } x_i \text{ and } x_j \text{ clusters with } x_k. \tag{1b}$$

Phase 3: Apply Bouwman's approach to form a causation tree for each cluster. Causation trees are defined in (2a) and (2b). These are similar to clusters except that only the root node of a causation tree must be on the list of symptoms. Any other element in the problem domain is a candidate for the causation tree for each cluster.

$$x_j \text{ causation tree } x_i \qquad \text{if } x_j \text{ impacts } x_i \text{ and } x_i \in S_t. \tag{2a}$$

$$x_k \text{ causation tree } x_i \qquad \text{if } x_k \text{ impacts } x_j \text{ and } x_j \text{ causation tree } x_i. \tag{2b}$$

Phase 4: Attempt to diagnose each cluster by constructing and testing arithmetic models based on the causation trees formed in Phase 3. A model is constructed for each x_i that is a root node in a causation tree. If the model accurately computes the observed change in the value of the root node, then a diagnosis may be obtained. If not, the model does not represent the problem domain faithfully and an accurate diagnosis is not possible.

Models are based on the structural modeling assumption that changes in x_i are due to changes in the variables that impact x_i. Thus we assume that

$$\Delta(x_i)=f(\Delta(x_j)\,|\,x_j \text{ impacts } x_i). \tag{3}$$

In order to compute $\Delta'(x_i)$, the estimate of $\Delta(x_i)$ given by the current structural model, we must have some information about the form of the function in (3). In the linear case, which will be used for simplicity of exposition, $\Delta'(x_i)$ is computed using

$$\Delta'(x_i)=\sum_j(C_{ji}\Delta(x_j)). \tag{4}$$

This is a breadth-first approach because all the variables one level below x_i enter into the analysis. It is a one-period pulse process.

If equation (4) is a perfect model for the actual changes in x_i, then no discrepancy will exist between the computed change $\Delta'(x_i)$ and the actual (observed) change $\Delta(x_i)$. Of course, in practice, discrepancies will occur. The accuracy of the model is reflected by the degree of discrepancy between the actual and the computed values. If the difference in actual and computed values is "small," then the model represents the problem domain faithfully and explains the change. That is, the model "explains" or accounts for the observed change, even if the change was not anticipated. (Recall that the diagnosis is attempting to explain unanticipated variances.)

As far as the system is capable of determining, the base causes of the problem lie along the branches of the causation tree. These branches represent different hypotheses about the cause of the problem. Of course, several branches may work together to contribute to the problem, while others may work to mitigate the problem somewhat. Branches contributing most to the problem (excitatory branches) will have an impact that is "large" relative to other branches and in the same direction (+ or −) as the change in the problem variable. Branches offsetting the problem by the greatest amount (inhibitory branches) have relatively large contributions, but in the opposite direction.

Branches causing and mitigating the problem constitute the first level of problem diagnosis. A deeper level of diagnosis is given by searching these branches for base causes, which lie at root nodes. If the causation tree models the problematical area of the domain accurately, then one or more variables at terminal nodes must be causing the problem. These variables have changed in value, thereby triggering changes up the branch to the root node or symptom variable. To discover these, branches of the causation tree are searched in an effort to find variables at terminal nodes that have changed and which lie along a branch whose impact is significant. These variables at terminal nodes constitute the ultimate diagnosis of the problem insofar as the system is capable of determining.

If the model is accurate and the system is able to find such nodes, then explanations such as "Variables x_1, x_2, and x_3 have contributed v_1, v_2, and v_3 to the changes observed in variable z; the problem would have been worse had it not been for variables y_1 and y_2 whose changes have offset the magnitude of the problem" can be offered. It also can provide explanations based on paths from terminal nodes to the root node.

Our discussion so far has assumed that the computed change $\Delta'(x_i)$ is "close to" the observed change $\Delta(x_i)$. If the difference in the computed change and the observed change is "large," the model fails to explain the source of the discrepancy. In this case, the model does not represent the structure of the problem domain accurately and the system is not able to diagnose the source of the problem. A message to that effect is delivered to the user. The process may be terminated at that point, or the model may be revised and retested. This is essentially a "what if. . ." type of process that helps the user explore and test hypothesized relationships in the problem domain. It is designed to help the user learn about the problem domain and formulate an accurate model of the structure of that domain. The next section describes how a history of successful diagnosis is maintained in the EKB and how this may be used to provide information concerning the causal cues described by Einhorn and Hogarth [41].

The EKB and Causal Cues

The EKB acts as the organizational memory for problem clusters that have been diagnosed successfully in the past, along with their respective diagnoses. It thus consists of structured declarative knowledge about problems and their diagnoses and can be represented in the format of a hierarchical data base (see Figure 4). Since a problem cluster may have occurred on several occasions, a count is maintained of the number of occurrences of each cluster.

In addition, clusters may have different diagnoses on different occasions, so a count also is maintained of the number of occurrences of each diagnosis, along with the corresponding dates and times. This information is useful in determining whether particular problem clusters have been occurring repeatedly, which may indicate some fundamental problem in organizational policy or strategy that should be addressed by management.

Figure 4: Logical structure of the EKB.

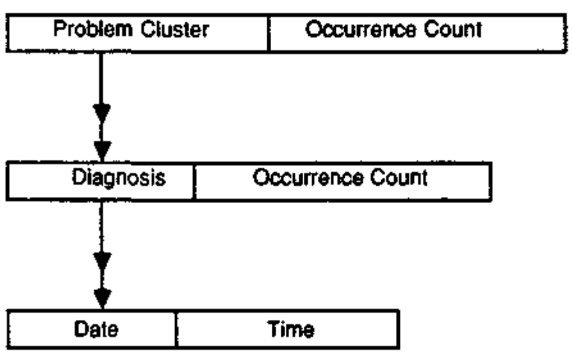

The EKB may be interpreted as simulating experiential human learning [54] [102] [57] [65]. In some cases, the EKB may be used to circumvent the need for a formal diagnosis. Whenever a cluster is formed, the EKB is searched first to determine whether the problem has been diagnosed previously. If so, the previous diagnoses are presented to the user, along with the current data values for variables constituting the diagnosis. If the user believes a previous diagnosis explains the current problem, the corresponding occurrence count is incremented and the current date/time stamp is inserted. If no previous diagnosis is acceptable, then a formal diagnosis is undertaken as described previously. If the problem is successfully diagnosed, appropriate action is taken to update the EKB.

The EKB, in conjunction with the diagnostic system, also is a potential source of information regarding the causal cues described by Einhorn and Hogarth [41] discussed earlier. The manner in which assertions are captured defines temporal order. The date/time stamp of each diagnosis clearly provides information about the contiguity in time of the cluster and its diagnosis. Since the EKB is for a single organization, contiguity in space is inherent in the system. Information about the number of alternative explanations is given by the number of different paths from terminal nodes to monitored variables. The EKB contains the number of times each of these paths has contributed to diagnoses that have occurred in the past. To generate information about the questions of whether Y would have occurred in the absence of X, the change in X can be set to zero in previous diagnoses and the system can be used to determine whether Y would have been on the list of problem symptoms. Finally, predictive ability (correlation coefficients) can easily be computed and similarity can be computed if weights are given for the cells.

We show elsewhere [88] [89] how causal modeling and path analysis, which are based on correlation and regression analysis, can be used to go beyond Einhorn and Hogarth's concept of predictive ability in helping users gain unbiased perceptions of relationships among variables in the problem domain.

APPLICATIONS

The system as described herein has been implemented and tested in prototype form by Ata Mohammed [8] as part of the BML/SLIM system used in other DSS research [32] [96] [62] [63]. Monitoring and generating the list of problem symptoms are based simply on changes in observed values from one period to the next and on user-supplied bounds for those values. The problem-structuring and diagnostic modules have been implemented in full, although linearity is assumed in the derived models and cycles are prohibited. The diagnostic process is under user control via menus, rather than being strictly phasewise.

This approach may appear to require very restrictive assumptions, but work already has begun to remove the linearity restriction and provide for automatic updating of impact coefficients by statistically deriving coefficients using data in the data base [88] [89].

Further, a large class of problems (diagnosis of accounting system data) satisfies the restrictions listed above. As illustrated in Figure 5, charts of accounts have an acyclic, hierarchical structure and relationships between levels are known and additive.

Moreover, accounting systems (which are one form of organizational model) are found universally in both profit and nonprofit organizations. Thus, even with very restrictive assumptions, the model is applicable to a type of problem common to *all* organizations.

Figure 5: An illustration of the hierarchical nature of accounting systems.

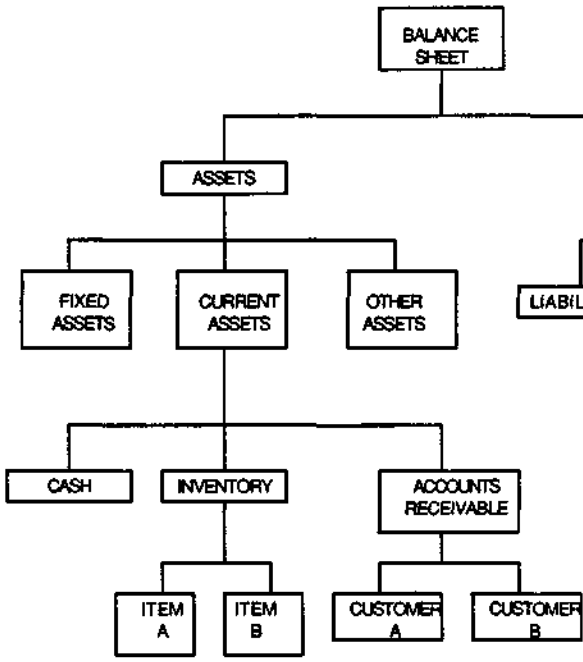

Moreover, it is standard practice in cost accounting systems to maintain standard costs against which actual costs are compared. In the budgeting process, most firms produce a master budget in which budgeted amounts are maintained for some accounts and against which actual costs are compared. Standard costs and budgeted amounts provide excellent vehicles on which the monitoring function and construction of causation trees may be based.

Large organizations maintain thousands of accounts in such systems. It may be exceedingly difficult for one individual to trace a large variance in current assets if the organization inventories thousands of items, has several thousand customers, or has several other current asset accounts. A diagnostic system such as the one we have described should greatly increase the speed and accuracy with which the causes of such variances can be isolated.

As an example, consider the recent emphasis on cost containment in the health care industry. The adoption of the Tax Equity and Fiscal Responsibility Act (TEFRA) of 1982 has caused a significant change in the way hospitals manage their daily operations. Prior to 1982, hospital staff members could order tests and prescribe medications and therapies with little regard for the cost incurred. This was because there frequently was a third party (e.g., Medicare) available to assume the responsibility for paying these costs.

The adoption of TEFRA has replaced the third-party cost reimbursement system with a cost-per-case payment system [80] [40] [14]. Put simply, each patient's case is classified into one of 467 classes known as diagnosis-related groups (DRGs). Third-party payment for a case is determined by the national average cost (adjusted for area wage differences) based on the case's DRG classification. If the hospital's costs for a particular case exceed the DRG payment, the hospital must absorb the difference. On the other hand, if the hospital's costs are lower than the DRG payment, the hospital enjoys a profit. The concept of ensuring that a profitable outcome is realized is known as "cost containment."

A system can be envisioned based on the methodology discussed above that could be used to support cost containment in hospital management. The decision maker would be the chief administrator in the hospital. The "problem" would be to identify processes (or services) that lead to inefficient handling of aspects of a patient's case.

Within a DRG many variables can be identified that contribute to the total cost of handling the case. For example, the length of the hospital stay is critical. Or the patient's age may necessitate special costly procedures. The number and type of drugs prescribed, the number of sutures and bandages required, and the equipment needed to administer these items also will impact the total cost. Nursing, laboratory analysis, and physical therapy services may be ordered. Finally, other costs such as heating, electricity, and administrative costs—known collectively as overhead costs—are incurred. A generic model of a DRG classification could be as shown in Figure 6.

The root node in the model in Figure 6 represents the payment amount for a particular DRG. This payment can be partitioned into five components (or submodels): special services, length of stay, material cost, service cost, and overhead cost. Each of these components can be partitioned again into finer subcomponents. For example, materials costs could be partitioned into costs for drugs and costs for other medical supplies. As can be seen in the model, this process can be carried out to any level of detail necessary.

By monitoring the variables in the model, the hospital's chief administrative officer would be able to identify areas in which substantial deviations caused hospital costs to exceed the DRG payment. This analysis could be restricted to specific DRG classifications—in which case the chief administrator would learn how well the hospital handles those particular cases—or the analysis could be across a range of DRG classifications. This would permit the administrator to examine how well the hospital functioned in a general area (e.g., in administering drugs). In this way, inefficient practices in the hospital's daily activities could be identified. Once identified, these activities could be studied and improved.

To illustrate, Averill, Kalison, Sparrow and Owens [10] suggested constructing departmental and physician "profiles" as a way of coping with the new era in health care management. A departmental profile based on our approach would examine

Figure 6: A generic DRG classification model.

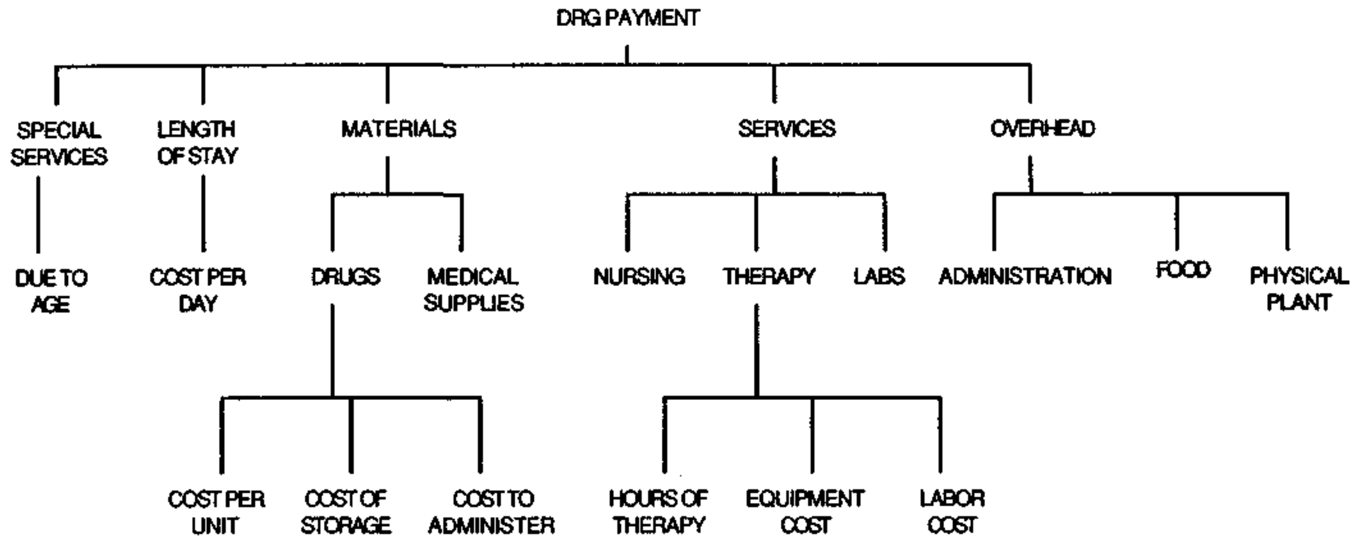

DRG models such as the one shown in Figure 6 to identify potentially inefficient resource usage. For example, suppose the costs for treating this DRG in an intensive care unit have increased 20 percent. The DRG model for the intensive care unit would identify where the cost increases had occurred.

A physician profile could be developed by examining all DRG models for a physician. By comparing a specific physician's profile against some designated "norm" (possibly constructed from mean or median values for all physicians), a physician's resource utilization pattern could be constructed. If a physician ordered an unusually large number of tests compared to the norm, for instance, a discussion with the physician to verify a need for these tests might be advised. The physician might be unaware of less expensive, equally effective methods of diagnosis.

Of course, this type of analysis could be extended in the opposite direction to hospital profiles. As the market share in the health care industry becomes more dominated by investor-owned "hospital chains" (e.g., Humana Corporation and Hospital Corporation of America), an entire hospital could be examined for overall efficiency relative to costs.

A system such as we describe would provide many immediate benefits to hospital administration. For example, storing a model for each of the 467 DRGs would be an imposing task for the mental capacity of most humans, but could be accomplished easily in a computer-based environment. (Since each patient case represents a unique instantiation of a model, some of the DRG models would be very complex and involve many variables. Even a small number of patient cases would generate a volume of data too great for efficient manual processing.) Most hospitals already maintain computer-based patient records so the marginal cost of data entry to provide a system such as we describe would be small. The payoff, on the other hand, would be tremendous. The system effectively would put the patient records to work for the hospital. Instead of passively generating billing statements, the records would take on an active role in hospital administration.

These example applications demonstrate two key points. First, the initial formulation in terms of linear models does not severely limit the immediate application of the methodology. On the contrary, linear accounting models are ubiquitous in business. Second, the methodology supports a learning process by the user. By building more intelligence into the software, the system could begin to share in the learning process. In the health care example, a system might be able to monitor the DRG case data automatically, deduce commonalities either within or across the DRG cases, and then suggest corrective measures whenever cost overruns appeared imminent. Since problem diagnosis is not unique to the health care industry, we believe this approach has potential for equal success in other industries as well.

SUMMARY AND FUTURE RESEARCH

To summarize, the system described in this paper supports aspects of the decision-making process from problem finding to problem structuring and diagnosis. When triggered by the monitor, the problem processor initiates a search for relevant assertions from the knowledge base (via the process control subsystem) and dynamically clusters these into causation trees representing potential explanatory problem hypotheses. To test the validity of these hypotheses, the system constructs an arithmetic representation of each branch and retrieves data from the data base to compute values based on the model. If the computed values are "close" to observed values, the user's mental model is validated and the problem's causes have been located. If the model fails to explain the problem, the user's model apparently is deficient and the system may be used to alter and retest the model. If a better model is found, the knowledge base is updated and used in future diagnoses. In this manner, the user and the system learn together about the underlying structure of the problem domain as models are tested and refined.

Our work extends that of Bouwman [23] in numerous ways: (1) it operationalizes his model within the context of ongoing managerial diagnostic activity and provides a system specifically designed for decision support; (2) it integrates Bouwman's approach with a data-base management system and provides for the construction of mathematical models to support ongoing quantitative diagnosis of problems; (3) it integrates Bouwman's approach with structural modeling and permits the maintenance of user-specified relationships in a knowledge base that can be modified as the user learns about the organizational problem domain; (4) it permits the use of digraph analysis techniques such as matrix powering, pulse analysis, and other techniques to support not only diagnosis but also prediction and analysis (including "what if. . ." analysis, the sine qua non of DSS modeling); (5) it includes a memory of successful diagnoses in the EKB, permitting the formal diagnostic process to be bypassed in certain circumstances; (6) it allows generation of information concerning causal cues via the EKB; and (7) as shown by Paradice and Courtney [88] [89], it permits direct extension to statistical techniques that may be used in numerous ways to support the acquisition, maintenance, and use of organizational knowledge and also permits the inclusion of noncausal relationships and certainty factors [90].

Future efforts will be directed to several activities. First, formal techniques of causal modeling and path analysis [7] will be incorporated into the system to test the validity of user-asserted relationships and to search for missing links in the model. Second, attempts will be made to incorporate other kinds of relationships in the knowledge base, such as "*A* is an upper bound for *B*," as well as strictly qualitative relationships. Efforts will be made to use these relationships to generate models automatically that will support subsequent phases of the decision process. Issues of model validation such as those described by Gass [50] and DeMillo, Lipton, and Perlis [38] may be addressed. Finally, laboratory studies may be conducted to determine the system's ability to influence cognitive biases such as those described by numerous researchers [31] [45] [85] [74] [109] [110] [111] [59] [60] [98] [2] [87] [108] [101] [12]. [Received: March 24, 1986. Accepted: January 28, 1987.]

REFERENCES

[1] Ackoff, R. L. Towards a system of systems concepts. *Management Science*, 1971, *17*, 661-671.

[2] Ackoff, R. L. The future of operational research is past. *Journal of the Operations Research Society*, 1979, *30*(2), 93-104.

[3] Ackoff, R. L. *Creating the corporate future.* New York: Wiley, 1981.
[4] Ackoff, R. L. Beyond prediction and preparation. *Journal of Management Studies,* 1983, *20*(1), 59-69.
[5] Ackoff, R. L., & Sasieni, M. W. *Fundamentals of operations research.* New York: Wiley, 1968.
[6] Adams, D., Courtney, J. F., & Kasper, G. M. Evaluation of DSS modeling software: A methodology and some results. In *Proceedings of the Twentieth Hawaii International Conference on Systems Sciences.* North Hollywood, CA: Western Periodicals, 1987.
[7] Asher, H. B. *Causal modeling.* Beverly Hills, CA: Sage, 1983.
[8] Ata Mohammed, N. H. *A knowledge based decision support system for managerial problem recognition and diagnosis.* Unpublished Ph.D. dissertation, Texas Tech University, 1985.
[9] Austin, L. M., & Burns, J. R. *Management science: An aid for managerial decision making.* New York: Macmillan, 1985.
[10] Averill, R. F., Kalison, M. J., Sparrow, D. A., & Owens, T. R. How hospital managers should respond to PPS. *Healthcare Financial Management,* 1984, *14*(3), 72-74, 76, 82, 84, 86.
[11] Axelrod, R. M. (Ed.). *The structure of decision.* Princeton, NJ: Princeton University Press, 1976.
[12] Barnes, J. H. Cognitive biases and their impact on strategic planning. *Strategic Management Journal,* 1984, *5*, 129-137.
[13] Basden, A. On the application of expert systems. *International Journal of Man-Machine Studies,* 1983, *19*, 461-477.
[14] Berry, L. E. A new treatment for health care cost. *Management Accounting*, 1984, *65*(10), 58-61.
[15] Blanning, R. W. Management applications of expert systems. *Information and Management,* 1984, *7*, 311-316.
[16] Blanning, R. W. An entity-relationship approach to model management. *Decision Support Systems,* 1986, *2*(1), 65-72.
[17] Blum, R. L. Discovery, confirmation, and incorporation of causal relationships from a large time-oriented clinical data base: The RX project. *Computer and Biomedical Research,* 1982, *15*, 164-187.
[18] Bonczek, R. H., Holsapple, C. W., & Whinston, A. B. Computer based support of organizational decision making. *Decision Sciences,* 1979, *10*, 268-291.
[19] Bonczek, R. H., Holsapple, C. W., & Whinston, A. B. The integration of network data base management and problem resolution. *Information Systems,* 1979, *4*(2), 143-154.
[20] Bonczek, R. H., Holsapple, C. W., & Whinston, A. B. The evolving roles of models in decision support systems. *Decision Sciences,* 1980, *11*, 337-356.
[21] Bonczek, R. H., Holsapple, C. W., & Whinston, A. B. Future directions for developing decision support systems. *Decision Sciences,* 1980, *11*, 616-631.
[22] Bonczek, R. H., Holsapple, C. W., & Whinston, A. B. *Foundations of decision support systems.* New York: Academic Press, 1981.
[23] Bouwman, M. J. Human diagnostic reasoning by computer: An illustration from financial analysis. *Management Science,* 1983, *29*, 653-672.
[24] Brown, A. Expert systems and PROLOG. *IEEE Videoconference Seminars via Satellite.* Produced by the Institute of Electrical and Electronics Engineers, Inc., and the Learning Channel, 1985.
[25] Brown, J. S., Burton, R., & deKleer, J. Pedagogical and knowledge engineering techniques in SOPHIE I, II, and III. In D. H. Sleeman & J. S. Brown (Eds.), *Intelligent tutoring systems.* London: Academic Press, 1982.
[26] Burns, J. R., & Winstead, W. H. An input/output approach to the structural analysis of digraphs. *IEEE Transactions on Systems, Man, and Cybernetics,* 1982, *SMC-12*, 15-24.
[27] Burns, J. R., & Winstead, W. H. M-labeled digraphs and Forrester-style causal models. *Management Science,* 1985, *31*, 343-357.
[28] Chandrasekaran, B. Towards a taxonomy of problem solving types. *AI Magazine,* 1983, *4*(1), 9-17.
[29] Chandrasekaran, B., & Mittal, S. Deep versus compiled knowledge approaches to diagnostic problem solving. *International Journal of Man-Machine Studies,* 1983, *19*, 425-436.
[30] Chorba, R. W., & New, J. L. Information support for decision maker learning in a competitive environment: An empirical study. *Decision Sciences*, 1980, *11*, 603-615.
[31] Churchman, C. W. *The design of inquiring systems.* New York: Basic Books, 1971.

[32] Courtney, J. F., DeSanctis, G. R., & Kasper, G. M. Continuity in MIS/DSS research: The case for a common gaming simulator. *Decision Sciences*, 1983, *14*, 419-439.
[33] Davis, R. Expert systems: Where do we go from here? *AI Magazine*, 1982, *3*(2), 3-22.
[34] Davis, R. Reasoning from first principles in electronic troubleshooting. *International Journal of Man-Machine Studies*, 1983, *19*, 403-423.
[35] Davis, R., Austin, H., Carlbom, I., Frawley, B., Pruchnik, P., Sneiderman, R., & Gilreath, J. The dipmeter advisor: Interpretation of geologic signals. In *Proceedings of the International Joint Conference on Artificial Intelligence*. Menlo Park, CA: American Association for Artificial Intelligence, 1981.
[36] Davis, R., Shrobe, H. E., Hamscher, W., Wickert, K., Shirley, M., & Polit, S. Diagnosis based on description of structure and function. In *Proceedings of the American Association for Artificial Intelligence*. Menlo Park, CA: AAAI, 1982.
[37] deKleer, J. A qualitative physics based on confluences. *Artificial Intelligence*, 1984, *24*(1-3), 7-83.
[38] DeMillo, R. A., Lipton, R. J., & Perlis, A. J. Social processes and proofs of theorems and programs. *Communications of the ACM*, 1979, *22*, 271-280.
[39] Dhar, V. On the plausibility and scope of expert systems in management. In *Proceedings of the Nineteenth Hawaii International Conference on Systems Sciences*. North Hollywood, CA: Western Periodicals, 1986.
[40] Doremus, H. D. A reimbursement system that limits the costs of hospital care. *Healthcare Financial Management*, 1983, *13*(4), 86-91.
[41] Einhorn, H. J., & Hogarth, R. M. Prediction, diagnosis and casual thinking in forecasting. *Journal of Forecasting*, 1982, *1*(1), 23-36.
[42] Elam, J. J. *Model management systems: A framework for development* (Technical Report 79-02-04). Unpublished manuscript, University of Pennsylvania, 1979.
[43] Elam, J. J., & Henderson, J. C. Knowledge engineering concepts for decision support system design and implementation. *Information & Management*, 1983, *6*(2), 109-114.
[44] Elam, J. J., Henderson, J. C., & Miller, L. W. Model management systems: An approach to decision support in complex organizations. In *Proceedings of the First International Conference on Information Systems*. Chicago, IL: Society for Management Information Systems, 1980.
[45] Feather, N. *Expectancy, incentive, and action*. Hillsdale, NJ: Lawrence Erlbaum, 1981.
[46] Fedorowicz, J., & Williams, G. D. Representing model knowledge in a logic-based DSS. *Decision Support Systems*, 1986, *2*(1), 3-14.
[47] Feigenbaum, E. A., & McCorduck, P. *The fifth generation: Artificial intelligence and Japan's computer challenge to the world*. Reading, MA: Addison-Wesley, 1983.
[48] Forbus, K. Modeling motion with qualitative process theory. In *Proceedings of the National Conference on Artificial Intelligence*. Menlo Park, CA: American Association for Artificial Intelligence, 1982.
[49] Gaschnig, J. Preliminary evaluation of the performance of the PROSPECTOR system for mineral exploration. In *Proceedings of the International Joint Conference on Artificial Intelligence*. Los Altos, CA: Morgan Kaufmann, 1979.
[50] Gass, S. I. Decision-aiding models: Validation, assessment, and related issues for policy analysis. *Operations Research*, 1983, *31*, 603-631.
[51] Gentner, D., & Stevens, A. L. (Eds.). *Mental models*. Hillsdale, NJ: Lawrence Erlbaum, 1983.
[52] Harary, F., Norman, R. Z., & Cartwright, D. *Structural models: An introduction to the theory of directed graphs*. New York: Wiley, 1965.
[53] Hart, P. E. Direction of AI in the eighties. *SIGART Newsletter*, 1982, No. 79, 11-16.
[54] Henderson, J. C., & Martinko, M. J. Cognitive learning theory and the design of decision support systems. *DSS-81 Transactions*. Atlanta, GA: Execucom Systems Corp., 1981.
[55] Hogarth, R. M., & Makridakis, S. Forecasting and planning: An evaluation. *Management Science*, 1981, *27*, 115-138.
[56] Hudetz, W. A graphical interactive system simulation program for minicomputers. In H. Bossel (Ed.), *Concepts and tools of computer-assisted policy analysis*. Basel, Switzerland: Birkhauser, 1977.
[57] Hunt, R. G., & Sanders, G. L. Propaedeutics of decision-making: Supporting managerial learning and innovation. *Decision Support Systems*, 1986, *2*(2), 125-134.

[58] Janis, I. L., & Mann, L. *Decision making: A psychological analysis of conflict, choice, and commitment.* New York: Free Press, 1977.
[59] Johnson, W. *People in quandries: The semantics of personal adjustment.* New York: Harper & Bros., 1946.
[60] Judson, A. J., & Cofer, C. N. Reasoning as an associative process: I. Direction in a simple verbal problem. *Psychological Reports,* 1956, *2,* 469-473.
[61] Kane, J. A primer for a new cross-impact language—KSIM. *Technological Forecasting and Social Change,* 1972, *4,* 129-142.
[62] Kasper, G. M. The effect of user-developed DSS applications on forecasting decision-making performance. *Journal of Management Information Systems,* 1985, *2*(2), 26-39.
[63] Kasper, G. M., & Cerveny, R. P. A laboratory study of user characteristics and decision-making performance in end-user computing. *Information and Management,* 1985, *9*(2), 87-96.
[64] Keen, P. G. W. MIS research: Reference disciplines and a cumulative tradition. In *Proceedings of the First International Conference on Information Systems.* Chicago, IL: Society for Management Information Systems, 1980.
[65] Kolb, D. A. *Applications of experiential learning theory to the information sciences* (working paper). Unpublished manuscript, Case Western Reserve University, 1978.
[66] Konsynski, B., & Dolk, D. Knowledge abstractions in model management. *DSS-82 Transactions.* Atlanta, GA: Execucom Systems Corp., 1982.
[67] Kruskal, J. B. Multidimensional scaling and other methods for discovering structure. In K. Enslein, A. J. Ralston, & H. S. Wilf (Eds.), *Statistical methods for digital computers.* New York: Wiley, 1977.
[68] Kruzick, P. G. *KSIM techniques for evaluating interaction among variables* (Technical note OED-16). Menlo Park, CA: Stanford Research Institute, 1973.
[69] Kuipers, B. Commonsense reasoning about causality: Deriving behavior from structure. *Artificial Intelligence,* 1984, *24*(1-3), 169-203.
[70] Kuipers, B. Qualitative simulation. *Artificial Intelligence,* 1986, *29,* 289-338.
[71] Langhorst, F. E. *Computer graphics aided interpretive structural modeling: A tool for conceptualizing complex design problems.* Unpublished Ph.D. dissertation, Purdue University, 1977.
[72] Lenard, M. L. Treating models as data. *Journal of Management Information Systems,* 1986, *2*(4), 36-48.
[73] Lendaris, G. G. Structural modeling—A tutorial guide. *IEEE Transactions on Systems, Man, and Cybernetics,* 1980, *SMC-10,* 807-840.
[74] Locks, M. O. The logic of policy as argument. *Management Science,* 1985, *31,* 109-114.
[75] Loy, S. L. *An experimental investigation of a graphical problem-structuring aid and nominal group technique for group decision support system.* Unpublished Ph.D. dissertation, Texas Tech University, 1986.
[76] McLean, J. M. Getting the problem right—A role for structural modeling. In H. A. Linstone & W. H. C. Simmonds (Eds.), *Future research: New directions.* Reading, MA: Addison-Wesley, 1977.
[77] McLean, J. M., & Shepherd, P. The importance of model structure. *Futures,* 1976, *18*(1), 40-51.
[78] McLean, J. M., Shepherd, P., & Curnow, R. *Techniques for analysis of system structure* (SPRU Occasional Paper Series No. 1). Sussex, England: University of Sussex, Science Policy Research Unit, 1976.
[79] Meador, C. L., Guyote, M. J., & Keen, P. G. W. Setting priorities for DSS development. *MIS Quarterly,* 1984, *8*(2), 117-129.
[80] Messmer, V. C. Methods that can be applied to DRG classifications. *Healthcare Financial Management,* 1984, *38*(1), 44-48.
[81] Michie, D. High road and low road programs. *AI Magazine,* 1982, *3*(1), 21-22.
[82] Minch, R. P., & Burns, J. R. Conceptual design for decision support systems utilizing management science models. *IEEE Transactions on Systems, Man and Cybernetics,* 1983, *SMC-13,* 549-557.
[83] Mintzberg, H., Raisinghani, D., & Théorêt, A. The structure of "unstructured" decision processes. *Administrative Science Quarterly,* 1976, *21,* 246-275.
[84] Mitroff, I. I., & Featheringham, T. R. On systemic problem solving and the error of the third kind. *Behavioral Science,* 1974, *19*(6), 383-393.

[85] Mitroff, I. I., Mason, R. O., & Barabba, V. P. Policy as argument—A logic for ill-structured decision problems. *Management Science,* 1982, *28,* 1391-1404.
[86] Mylopoulos, J., Shibahara, T., & Tsotsos, J. Building knowledge-based systems: The PSN experience. *IEEE Computer,* 1983, *16*(10), 83-89.
[87] Newell, A., & Simon, H. A. *Human problem solving.* Englewood Cliffs, NJ: Prentice-Hall, 1972.
[88] Paradice, D. B., & Courtney, J. F. Controlling bias in user assertions in expert decision support systems for problem formulation. *Journal of Management Information Systems,* 1986, *3*(1), 52-64.
[89] Paradice, D. B., & Courtney, J. F. SmartSLIM: A DSS for controlling biases during problem formulation. In *Proceedings of the TAMU Human Factors in MIS Conference.* Norwood, NJ: Ablex, 1986.
[90] Paradice, D. B., & Courtney, J. F. Causal and non-causal relationships and dynamic model construction in a managerial advisory system. *Journal of Management Information Systems,* 3(4), 39-53.
[91] Patil, R., Szolovits, P., & Schwartz, W. Causal understanding of patient illness in medical diagnosis. In *Proceedings of the International Joint Conference on Artificial Intelligence.* Menlo Park, CA: American Association for Artificial Intelligence, 1981.
[92] Pople, H. Heuristic methods for imposing structure on ill-structured problems. In I. DeLotto & M. Stefanelli (Eds.), *Artificial intelligence in medicine.* New York: North Holland, 1985.
[93] Pounds, W. F. The process of problem finding. *Industrial Management Review,* 1969, *11*(1), 1-19.
[94] Pracht, W. E. *An experimental investigation of a graphical interactive structural modeling aid for decision support systems.* Unpublished Ph.D. dissertation, Texas Tech University, 1984.
[95] Pracht, W. E. GISMO: A visual problem-structuring and knowledge-organization tool. *IEEE Transactions on Systems, Man and Cybernetics,* 1986, *16,* 265-270.
[96] Pracht, W. E., & Courtney, J. F. *The effects of an interactive graphics based DSS to support problem structuring* (working paper). Unpublished manuscript, Texas A&M University, Business Analysis and Research Department, 1986.
[97] Raiffa, H. *Decision analysis.* Reading, MA: Addison-Wesley, 1968.
[98] Reitman, W. R. Heuristic decision procedures, open constraints, and the structure of ill-defined problems. In M. W. Shelly II & G. L. Bryan (Eds.), *Human judgments and optimality.* New York: Wiley, 1964.
[99] Roberts, F. S. Building and analyzing an energy demand signed digraph. *Environment and Planning,* 1973, *5*(3), 18-29.
[100] Roberts, F. S. *Graph theory and its applications to problems of society.* Philadephia, PA: Society for Industrial and Applied Mathematics, 1978.
[101] Rowe, A. J. How do senior managers make decisions? *Business and Economics,* Winter 1977, pp. 17-20.
[102] Sage, A. P. Behavioral and organizational considerations in the design of information systems and processes for planning and decision support. *IEEE Transactions on Systems, Man, and Cybernetics,* 1981, *SMC-11,* 640-678.
[103] Schueller, J. K., Slusher, R. M., & Morgan, S. M. An expert system with speech synthesis for troubleshooting grain combine performance. *Transactions of the American Society of Agricultural Engineers,* 1986, *29,* 342-344.
[104] Sen, A., & Biswas, G. Decision support systems: An expert systems approach. *Decision Support Systems,* 1985, *1*(3), 197-204.
[105] Shortliffe, E. *Computer-based medical consultations: MYCIN.* New York: American Elsevier, 1976.
[106] Simon, H. A. *The new science of management decisions.* New York: Harper & Row, 1960.
[107] Stefik, M., Aikens, J., Balzer, R., Benoit, J., Birnbaum, L., Hayes-Roth, F., & Sacerdoti, E. The organization of expert systems, A tutorial. *Artificial Intelligence,* 1982, *18,* 135-174.
[108] Taylor, I. A. A retrospective view of creativity investigation. In I. A. Taylor & J. W. Getzels (Eds.), *Perspectives in creativity.* Chicago, IL: Aldine, 1975.
[109] Tversky, A., & Kahneman, D. Belief in the law of small numbers. *Psychological Bulletin,* 1971, *76*(2), 105-110.
[110] Tversky, A., & Kahneman, D. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology,* 1973, *5,* 207-232.

[111] Tversky, A., & Kahneman, D. Judgement under uncertainty: Heuristics and biases. *Science*, 1974, *185*, 1124-1131.
[112] Tversky, A., & Kahneman, D. The framing of decisions and the psychology of choice. *Science*, 1981, *211*, 453-458.
[113] Wakeland, W. QSIM2: A low budget heuristic approach to modeling and forecasting. *Technological Forecasting and Social Change*, 1976, *9*, 213-229.
[114] Waller, R. J. Comparing and combining structural models of complex systems. *IEEE Transactions on Systems, Man and Cybernetics*, 1979, *SMC-9*, 1120-1128.
[115] Wang, M. S., & Courtney, J. F. Design and implementation of the MAGIC/ROC decision support systems generator. *DSS-82 Transactions*. Atlanta, GA: Execucom Systems Corp., 1982.
[116] Wang, M. S., & Courtney, J. F. A conceptual architecture for generalized DSS software. *IEEE Transactions on Systems, Man and Cybernetics*, 1984, *24*, 701-710.
[117] Warfield, J. N. *Structuring complex systems* (Monograph No. 4). Columbus, OH: Battelle Memorial Institute, 1974.
[118] Warfield, J. N. *Societal systems: Planning, policy and complexity*. New York: Wiley-Interscience, 1976.
[119] Weiss, S. M., Kulikowski, C. A., Amarel, S., & Safir, S. A model-based method for computer-aided medical decision making. *Artificial Intelligence*, 1978, *11*(2), 145-172.
[120] Yadav, S., & Korukonda, A. Management of Type III error in problem identification. *Interfaces*, 1985, *15*(4), 55-61.

Jim Courtney holds an endowed professorship in the Business Analysis and Research Department at Texas A&M University. He received his Ph.D. in business administration (management science) from the University of Texas at Austin. Besides his other academic experiences, Dr. Courtney has held positions as Database Analyst at MRI Systems Corporation (presently INTEL Corporation's Austin division) and Visiting Research Scientist at NASA's Johnson Space Center. His papers have appeared in several journals including *Management Science, Communications of the ACM, IEEE Transactions on Systems, Man and Cybernetics, MIS Quarterly, Decision Sciences,* the *Journal of Management Information Systems, Database, Interfaces,* the *Journal of Applied Systems Analysis,* the *Journal of Bank Research, Socio-Economic Planning Sciences,* and the *Journal of Experiential Learning and Simulation.* He is codeveloper of the Systems Laboratory for Information Management, a software package to support research and education in decision support systems. His present research interests are in knowledge-based decision support systems, data-base systems, and management information systems.

David Paradice is Assistant Professor of MIS in the Department of Business Analysis and Research at Texas A&M University. He holds a B.S. in information and computer science and an M.S. in industrial management from Georgia Tech and a Ph.D. in business administration (MIS) from Texas Tech. Dr. Paradice's primary interests are in investigating means for applying artificial intelligence techniques to managerial information systems. Related research includes investigating methods of dynamic model construction in managerial advisory systems and developing methods of utilizing qualitative models of organizational systems. Aspects of his research have appeared in the *Journal of MIS.*

Nassar H. Ata Mohammed is on the faculty of the King Faud University of Petroleum, Dhahran, Saudi Arabia. He holds B.S. and M.S. degrees in chemical engineering from the University of Texas at Austin and received a Ph.D. in business administration from Texas Tech University. Dr. Ata Mohammed previously was manager of computer consulting services at the King Faud University of Petroleum and has worked as an information systems consultant for several organizations in Saudi Arabia and the United States. He was a speaker at the Third National Computer Conference in Jeddah. His research interests are in knowledge-based decision support systems.