
Causal and Non-causal Relationships and Dynamic Model Construction in a Managerial Advisory System

DAVID B. PARADICE and JAMES F. COURTNEY, JR.

DAVID PARADICE is an assistant professor in the Department of Business Analysis and Research at Texas A&M University. His research interests focus on managerial expert systems. Results from his prior research have appeared in the *Journal of Management Information Systems*.

JAMES F. COURTNEY, JR., holds an endowed professorship in Management Information Systems in the College of Business Administration at Texas A&M University. He received his Ph.D. in Business Administration (Management Science) from the University of Texas at Austin in 1974. His experience includes positions as Database Analyst with MRI Systems Corporation, Visiting Research Scientist at the National Aeronautics and Space Administration Johnson Space Center in Houston, Assistant Professor of Management Science at Georgia Tech, Project Director at Georgia Tech's Engineering Experiment Station, Visiting Professor at the State University of New York at Buffalo, and Professor of Management Information Systems at Texas Tech University. His papers have appeared in several journals, including *Management Science*, *Communications of the ACM*, *MIS Quarterly*, the *Journal of Management Information Systems*, *IEEE Transactions on Systems, Man, and Cybernetics*, *Database*, *Socio-Economic Planning Sciences*, *Applied Systems Analysis*, *Journal of Bank Research*, *Decision Sciences*, *Interfaces*, *Information and Management*, and the *Journal of Experiential Learning and Simulation*. He is co-developer of the Systems Laboratory for Information Management (Business Publications, Inc., 1981), a software system to support research and education in management information systems. His present research interests are in decision support systems, expert systems, and database management.

ABSTRACT: Recent successes in the design and implementation of expert systems in non-managerial problem domains has spurred interest in the proposition that expert systems may be viable in managerial problem domains as well. Continuing research in the design, implementation, and testing of a prototype managerial expert system is described in this paper. Issues regarding managerial knowledge organization, acquisition, representation, and utilization are discussed, and progress in the

This paper was presented at the Twentieth Annual Hawaii International Conference on System Sciences, Honolulu, January 6-9, 1987, and is published here with permission.

development of an advisory module is presented. Finally, results of evaluating the prototype versus the abilities of human subjects are also presented.

KEY WORDS AND PHRASES: Problem formulation, expert systems, decision support systems.

1. Introduction

O'KEEFE STATES in reference [7] that the development of expert systems (ES) consists of three major stages: (1) knowledge acquisition, (2) knowledge organization, and (3) knowledge utilization. We would add a fourth stage, knowledge maintenance, and a loop back to knowledge utilization. Previous papers have shown how graphical structural modeling techniques may be useful during knowledge acquisition and organization [10] and how causal modeling may be used to control user biases during the knowledge acquisition process [8].

This paper shows how the causal modeling approach has been extended to the knowledge organization, utilization, and maintenance phases of the ES development process. As before, one form of acquired knowledge consists of user-asserted causal relationships between a dependent variable and a set of independent variables in an organizational database. At the time of acquisition, values from the database are input into causal modeling procedures in an attempt to statistically validate the asserted causal relationships. Results of the causal modeling analysis are presented to the user, who may either reject relationships or have them inserted into the knowledge base. Rejected relationships are maintained in a "rejection base" for later use.

Belatedly, the previous system has been tabbed SmartSLIM I to distinguish it from the version described in this paper: SmartSLIM II. The name SmartSLIM is used because the system is an intelligent extension to the SLIM decision support system (DSS) package described elsewhere [5].

Extensions of the previous system are numerous. First, a taxonomy of non-causal relationships has been developed and is used as the basis for knowledge organization. A second extension is an advisory module capable of using causal knowledge to dynamically construct models in response to a user's request for information about a target variable. The system-constructed model and information about non-causal relationships are used to provide advice as to what variables may be manipulated to influence the target variable.

A third extension is the development of a "discovery" module, which constitutes a form of system learning. The discovery module uses information in the rejection base and data dictionary to construct possibly new causal models which are tested as before. Statistically valid models are inserted into the knowledge base with a flag indicating they were discovered by the system and have not been user-validated. They may be user-validated later and used in future advisory processes. A separate paper is being developed to present our work in this area.

The next section briefly reviews the conceptual basis and features of SmartSLIM

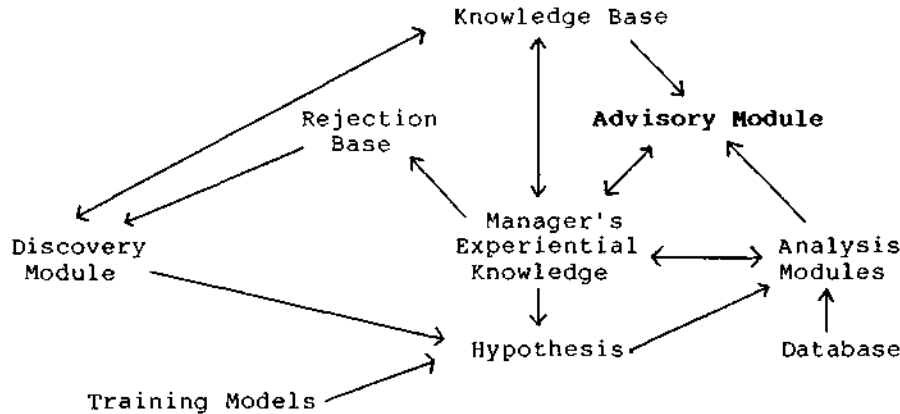


Figure 1. Conceptual Model of the SmartSLIM System

I. Then the taxonomy of relationships developed for acquiring and organizing knowledge is presented. Next, the unique features of SmartSLIM II are described. Finally, we describe how the system was evaluated by comparing SmartSLIM II's output on questionnaires about a simulated management domain to responses of human subjects to the same questionnaires.

2. The Conceptual Model and SmartSLIM I

FIGURE 1 SHOWS the conceptual model of the SmartSLIM I system. Both systems SmartSLIM I and II employ causal modeling techniques as a method of inferring knowledge about a business problem domain. Causal modeling is a statistical method particularly attractive as a decision-making aid in a business problem context, since causal relationships are frequently sought in problems in the business domain [2].

One form of causal modeling draws on the method of path analysis to examine relationships among variables. The goal of path analysis is to provide plausible explanations of observed correlations by constructing cause-and-effect models of relations among variables. Although an observed correlation can never be used as proof of a causal relationship, very convincing arguments for causality can be constructed from (1) statistical inference, (2) postulated relationships developed from knowledge of the subject matter, and (3) common sense.

Through repeated sessions with the SmartSLIM system, users hypothesize and test relationships between items in a simulated organizational database created by the Business Management Laboratory (BML) [6]. The database is cross-sectional in nature, containing data on over thirty firms in a single period in time. Hypotheses are tested by constructing linear (stepwise) regression models.

The results of the regression are presented to the user, at which time the user may




Figure 2. A Simple Non-reciprocal Causal Model

either accept or reject these results. On acceptance, the system asks for additional information regarding the nature of the relationship between the variables in the regression model and updates the knowledge base. On rejection, the system updates the rejection base. SmartSLIM I draws upon a "tool-kit" of structural modeling, path analysis, and linear statistical techniques to synthesize the knowledge in the knowledge base.

2.1 Knowledge Base Creation and Design

The CAUSE command is used in both SmartSLIM I and II to specify the variables which will compose a causal model. These models represent hypothesized relationships between variables. An example of the command is:

```
CAUSE SV(F1-8,Q1), PRICE(F1-8,Q1), ADV(F1-8,Q1);
```

In this example, SV is the database item "sales volume," PRICE is "price," and ADV is "advertising expense." The first variable in the variable list represents the endogenous variable in the analysis. All other variables are taken to be exogenous variables. Hence, this example postulates that price and advertising expense affect sales volume. This command tests the causal model shown in Figure 2.

Hypothesis tests which determine the significance of the coefficient estimates are interpreted by the software. Next, the user is asked whether the relationship should be retained for future reference. If the user responds affirmatively, the system begins prompting for other information relevant to this relationship including the type of relationship (discussed in Section 3), an indication of the user's agreement with the statistical results, an estimate of any time lag involved, up to three conditions affecting the relationship, and a qualitative description of the relationship.

2.2 Knowledge Base Perusal and Modification

Several commands were implemented in SmartSLIM I to facilitate inspection and modification of the knowledge base. The knowledge base is retained in main memory while the system is operating and is updated as the session ends. This approach has the advantage of extremely fast access to the contents of the knowledge base, but suffers the disadvantage of being particularly vulnerable to any type of system-related failure. The KBSAVE command allows the user to periodically save the contents of memory without leaving the system.

The RULES command is used to examine the statements that have been stored in the knowledge base describing the relationships, or qualitative output. The KBPRINT command generates predominately quantitative output captured during hypothesis testing. Modifications to the knowledge base are made via the KBMODIFY and the KBDELETE commands.

2.3 Knowledge Base Manipulation in SmartSLIM I

The knowledge base of SmartSLIM I is manipulated via path analysis. Path analysis is initiated by executing a PATH command, which utilizes a depth-first backtracking algorithm to create paths from the pairwise relationships in the knowledge base. This command searches for existence of indirect relationships between variables.

Cyclical paths are not considered by the PATH command. The command format is described in detail shortly. Briefly, the user must specify at least one variable, which serves to prune the possible paths that can be constructed. Only causal and compositional relationships are considered for inclusion in a path. This restriction maintains consistency with the MODEL command discussed below.

The PATH command is followed by either one database item name and the keyword ALL, or two database item names. Examples of the three general forms of the command are:

- (1) PATH ALL, SALES;
- (2) PATH PRICE, ALL;
- (3) PATH PRICE, SALES;

If ALL is specified before a database item name, then SmartSLIM I creates all paths which exist from any database item in the knowledge base to the database item specified after the ALL keyword. In other words, the system is finding all paths to a particular database item. This is analogous to finding all database items that affect the one specified. In the first example above, the system will construct a path for all variables that can be shown to ultimately impact the SALES variable (i.e., sales revenue).

If the database item name is specified before the ALL keyword, then all paths are constructed that begin with the specified database item. This is analogous to deter-

cost-decrease impacts (profits)
 advertising impacts (sales costs market-share)
 price-decrease impacts (sales market-share)
 producing-more impacts (inventory costs)
 better-quality impacts (sales)
 profit impacts (dividends bonuses)
 sales impacts (revenue profit)
 dividends impacts (stock-rating)
 X member-of (X : Y)
 X member-of (Y : Z) if X member-of Z
 X affects Y if X impacts Z and Y member-of Z
 (X Y) is-path if X affects Y
 (X (Y (Z))) is-path if X affects Y and (Y Z) is-path

Figure 3. A MicroProlog Implementation of the Path Construction Routines

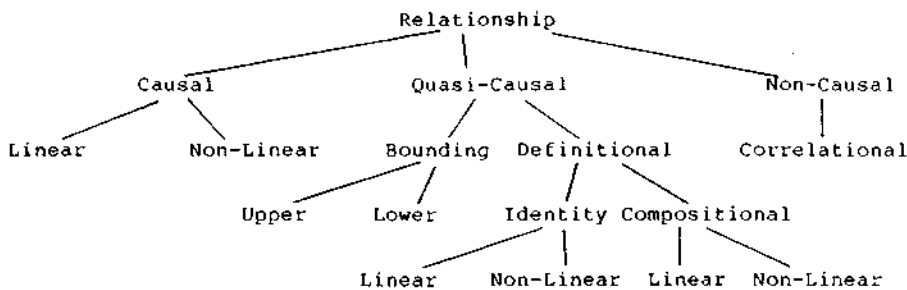


Figure 4. Taxonomy of Relationships in SmartSLIM

mining what variables are affected by the specified database item. In the second example above, the system will construct all paths from the PRICE variable to any other variable in the system.

When two database item names are specified, then only paths which begin with the first database item specified and end with the second are examined. In the third example above, the system will construct all paths from PRICE to SALES. Figure 3 contains a MicroProlog rule base showing how a general approach to a production rule implementation of the path construction could be realized.

3. A Taxonomy of Variable Relationships

EXPERIENCE IN USING SmartSLIM I to model relationships in BML indicates that causal knowledge alone is insufficient to model relationships in a business domain. This section describes a taxonomy of relationships which begins the process of classifying relationships more carefully so that they can be accurately represented in a knowledge base.

Variables in a business domain may interact in many ways. Based on our experience, it is useful to organize relationships into three broad categories: causal, quasi-causal, and non-causal (see Figure 4).

Causal relationships exist when the occurrence of a change in one variable affects a

change in a second variable in a theoretically justifiable manner. Non-causal relationships exist when two variables are correlated, that is, mathematically related without a theoretically justifiable causal implication. Quasi-causal relationships are causal under certain conditions (bounding) or are definitional in a mathematical sense.

Bounding relationships are lower and upper limiting constraints. A bounding variable may exhibit a causal influence if a second variable must exceed a specified value determined by the bounding variable. "Price" typically exceeds "cost of production," for instance. Similarly an upper limit may influence (restrictively) a variable. For example, "production" may be constrained by "capacity" or "raw material inventory."

Definitional relationships are categorized as either identity or compositional. An identity is a restatement of a variable in a new unit of measurement. "Sales volume" (measured in number of units sold) and "sales revenue" (measured in dollars) exemplify an identity relationship. A compositional relationship exists when a variable exists as a component of a second variable. For example, "depreciation" is a component of the calculation of "net income." A change in a variable in a quasi-causal relationship may "cause" a change in a second variable which is not accurately described as causal or correlational.

4. SmartSLIM II

DURING KNOWLEDGE ACQUISITION with the CAUSE command in SmartSLIM II, the user is asked to classify each relationship into one of four categories: causal, definitional, correlational, or compositional. A fifth category, bounding, is handled separately since one variable may both "cause" and "bound" a second. For example, a specified level of inventory causes a minimal (i.e., lower bound) of inventory charges to be incurred.

4.1 Modifications to the CAUSE and PATH Commands

Taxonomic information on relationships is stored in the knowledge base along with statistical data generated by the causal modeling analysis and important system-related data used to access the database. In cases where a relationship is not retained in the knowledge base, it is stored in the rejection base. The rejection base maintains data indicating which relationships have been rejected and how often.

The rejection base serves as a reservoir of relationships for which there is currently no statistical basis, but which may be fruitful areas for future investigation. The rejection base is used in later sessions to guide automatic discovery of relationships. The rejection base may also be examined for biases demonstrated by the user, since it retains how many times particular relationships have been examined.

The CAUSE command in SmartSLIM II extends to polynomial relationships. However, since space in this paper is limited and the facilities that manipulate the contents of the knowledge base are predominately restricted to linear relationships, methods for addressing polynomial relationships will not be discussed.

The PATH command has been extended to accommodate the new taxonomy of

relationships. Only causal and compositional relationships are considered for inclusion in a path. This restriction maintains consistency with the MODEL command discussed below.

In the case where both variables have been specified, the software performs an analysis of the paths as demonstrated in Asher [1] and Burns and Winstead [3, 4]. Detailed analysis of a path seeks to resolve two questions. The first question concerns whether indirect paths from the first variable to the last variable in the selected path indicate the need to include a direct path between these variables. The second question concerns whether the overall effect of all the indirect paths is excitatory, inhibitory, or neutral.

The first question is resolved by following Asher [1]. Asher shows that if all indirect paths which exist have been included in the knowledge base, then the sum of the products of the path coefficients for the indirect paths should equal the standardized regression coefficient estimate of the first variable regressed on the last (the correlation). Mathematically, this can be written:

$$r_{fl} = \sum_{i=1}^k \prod_{j=1}^{m_i} p_{ij}$$

where:

- f denotes the first variable in the paths,
- l denotes the last variable in the paths,
- k is the number of paths from f to l ,
- m_i is the number of arcs in path i ,
- p_{ij} is j -th coefficient in the i -th path, and
- r_{fl} is the correlation between f and l .

The sum of the products of the path coefficients is known as the "contribution" of the path. A statistical test of the hypothesis that the correlation is equal to the sum of the path contributions is being implemented. Currently, the two numbers are presented for the user to determine if they are "close enough" to be considered equal.

If the contribution and correlation values are close, this indicates that the indirect paths capture as much "information" as would be contained in a direct path from the first variable to the last variable. This being the case, there is no need to incorporate the direct relationship between the first and last variable into the knowledge base. If these two values are not approximately equal, then the user is advised that incorporating the direct relationship into the knowledge base via the CAUSE command should be considered.

The second question is resolved by taking the approach outlined by Burns and Winstead [3, 4], who show that paths in a structural (i.e., causal) model may be classified as either excitatory or inhibitory. An excitatory path has the characteristic that increases in the first variable in the path result in increases in the final variable. A path is excitatory if all the arcs in the path have positive coefficients, or if there is an even number of arcs with negative path coefficients. Inhibitory paths are characterized by increases in the first variable in the path resulting in decreases in the final variable. Inhibitory paths have an odd number of negative path coefficients.

Once the paths between two variables have been classified as either excitatory or inhibitory, the number of excitatory paths and the number of inhibitory paths can be plotted as a point in an excitatory-inhibitory plane. Subsequently, the angle of a line from the origin of the two axes defining the plane to the point can be computed. The value of the angle Θ defined line and the excitatory axis (X-axis) is given by:

$$\Theta = \cos^{-1} \left(\frac{E_{fl}}{\sqrt{(E_{fl}^2 + I_{fl}^2)}} \right)$$

where:

f denotes the first variable in the paths,

l denotes the last variable in the paths,

E_{fl} is the number of excitatory paths between f and l , and

I_{fl} is the number of inhibitory paths between f and l .

The position of the line in the plane gives an indication of the overall structural effect.

A final aspect of the analysis is SmartSLIM II's determination of the level of confidence it has in its calculations. The system considers the paths as redundant ways of achieving a change in one variable given a change in another. Therefore, the "weakest" link in the paths is the one which will be most suspect. The strength of the links is taken as the level of agreement specified by the user via the CAUSE command.

However, since the paths are redundant, the system takes an optimistic view of its model of the world and chooses the largest of the weakest links as its level of agreement with its calculations (i.e., a maxi-min approach). For example, if there are three indirect paths between two variables which have weakest links (minimum agreement levels) of 3, 4, and 7, then the system chooses a value of 7 for its belief in its final determination of the effect of the first variable in the paths on the last.

4.2 Construction of Models and the Advisory Module

One of the most powerful features of SmartSLIM II is the dynamic construction of models by the system itself. Model construction is invoked via the MODEL command. The format of the MODEL command is:

MODEL [INCREASE | DECREASE] (target variable);

For example, the command MODEL INCREASE SV; (as in Figure 5) means to construct a model with the objective of increasing the target variable sales volume (SV).

This command initiates a search of the knowledge base for all relationships in which the target variable participates. Definitional and correlational relationships are output first for user inspection. Then the system constructs a model using causal and compositional relationships, and estimates the coefficients of the model.

> MODEL INCREASE SV;

THE FOLLOWING IS KNOWN ABOUT SV:

SV IS CORRELATED WITH RMI.
SV IS CORRELATED WITH PVFS.
SV IS DEFINED BY SALE.

THE FOLLOWING MODEL HAS BEEN CONSTRUCTED:

CAU SV(F1-*) ,ADV(F1-*) ,PRICE(F1-*) ,SREP(F1-*) ,
BO(F1-*) ,QC(F1-*) ;

ESTIMATES FOR THE MODEL ARE:

0.6824 ADV -0.8145 PRICE -0.3702 BO
BO HAS A LAGGED EFFECT NOT WELL MODELED HERE.

SREP WAS EXCLUDED FROM THE MODEL.

QC WAS EXCLUDED FROM THE MODEL.

THIS MODEL EXPLAINS 82.5881% OF THE VARIATION IN SV.

SV HAS THE FOLLOWING POLYNOMIAL RELATIONSHIP WITH ADV:

$SV = 0.1822 E06 + 0.90580 E05 ADV - 0.55090 E05 ADV ^ 2$

IN ORDER TO INCREASE SV

INCREASE ADV
DECREASE PRICE
INCREASE RMI
DECREASE BO

TYPE "MORE" FOR MORE MODELING INFORMATION,
TYPE "END" TO EXIT THE MODEL COMMAND,
OR TYPE ONE OF THE VARIABLE NAMES JUST LISTED
FOR MORE INFORMATION ABOUT THAT RELATIONSHIP: PRICE

A ONE UNIT CHANGE IN PRICE PRODUCES A -0.638 DECREASE
IN SV. AGREEMENT LEVEL FOR THIS RELATIONSHIP IS 8.

TYPE "MORE" FOR MORE MODELING INFORMATION,
TYPE "END" TO EXIT THE MODEL COMMAND,
OR TYPE ONE OF THE VARIABLE NAMES JUST LISTED
FOR MORE INFORMATION ABOUT THAT RELATIONSHIP: MORE

IN ORDER TO DECREASE PRICE
DECREASE COM
DECREASE ADV
DECREASE QC

TYPE "MORE" FOR MORE MODELING INFORMATION,
TYPE "END" TO EXIT THE MODEL COMMAND,
OR TYPE ONE OF THE VARIABLE NAMES JUST LISTED
FOR MORE INFORMATION ABOUT THAT RELATIONSHIP: END

Figure 5. Advisory Action of System regarding Goal of Increasing Sales Volume

Only causal and compositional relationships are used to construct the model in order to maintain the stability of the resulting model. Definitional and correlational relationships are highly correlated. Consequently, inclusion of these variables in the models would lead to coefficient estimates with large variances. A model constructed in this fashion (i.e., including highly correlated exogenous variables) characteristically has a high R^2 , but no significant coefficient estimates.

A previously untested model could be developed at this stage, since the system consolidates *all* its knowledge regarding the database item of interest. For example, this could result in consolidated relationships with a common variable acquired from different variables.

If the model is a linear model, then a stepwise regression approach is used. If the model is polynomial, a least squares approach is used. Where lagged variables exist in the linear models, a note is made of their presence, but the code for lagged parameter estimation has not yet been incorporated into the system.

When the optional modifiers INCREASE or DECREASE are specified, the system provides a list of variables that may be manipulated in order to achieve the desired result. For example, if the user is interested in increasing the value of a variable, the system indicates which variables may be altered, and in what direction, in order to achieve this goal. In this way, the system assumes an advisory role, effectively offering alternative courses of action to be considered. In cases where conditional information exists, this information is also presented.

A second feature of the INCREASE or DECREASE option is that the system allows a user to examine variables that are introduced in the process of constructing a model. If, for instance, the system suggests decreasing price in order to achieve a goal of increasing sales volume, the user can immediately solicit information regarding a new goal of decreasing price. This process can be repeated as often as desired.

5. Evaluation of the Methodology

SmartSLIM HAS BEEN EVALUATED using a test given human subjects. In research by Pracht [9], subjects were given a post-experiment questionnaire which asked (1) the level of confidence the subjects felt in making certain BML decisions, (2) the degree of structure the subjects perceived in certain BML decision-making situations, and (3) the relationships the subjects believed existed in the BML simulation.

The PATH and MODEL commands have been used to answer the first two parts of the questionnaire developed by Pracht. Tables 1, 2, and 3 were constructed to aid in evaluating the system. Pracht's questionnaire asked the subjects to rate their confidence in the accuracy of certain decision contexts. The subjects were given a five-point scale labeled "not at all confident" to "highly confident." (The system generates confidence levels on a ten-point scale.)

On examining the tables, we see that the system exhibited high confidence in its ability to model some areas of the marketing decision. Confidence in the areas of plant and production decisions was generally lower. In comparison, the mode for the

Table 1 Marketing Decisions Affecting Sales Volume (results of path analysis of variables affecting marketing decisions)

Variables	Correlation	Contribution	Confidence	Effect
Price	-0.446	N/A	8	Inhibitory
Sales representatives	-0.549	N/A	7	Inhibitory
Advertising	0.602	0.383	7	Neutral
Commission	-0.178	-0.135	5	Inhibitory
Quality control	0.478	-0.283	7	Inhibitory

Note: N/A = not applicable.

Table 2 Plant and Production Decisions Affecting Sales Volume (results of path analysis of variables affecting production decision regarding sales volume)

Variables	Correlation	Contribution	Confidence	Effect
Raw material orders	0.894	0.232	5	Neutral
Production volume	0.953	0.378	5	Neutral
Maintenance expense	0.756	-0.028	7	Inhibitory
Plant capacity	0.832	0.510	7	Excitatory

Table 3 Plant and Production Decisions Affecting Cost of Goods Sold (results of path analysis of variables affecting production decisions regarding cost of goods sold)

Variables	Correlation	Contribution	Confidence	Effect
Raw material orders	0.874	N/A	4	Excitatory
Production volume	0.949	0.455	3	Excitatory
Maintenance expense	0.707	N/A	5	Excitatory
Plant capacity	0.805	N/A	5	Excitatory

Note: N/A = not applicable.

responses to each of these decision contexts given by the subjects in Pracht's study was "somewhat confident," or a 4 on a scale from 1 to 5.

The system was totally unable to address decision contexts involving finance and administration, a third area tested in the first part of Pracht's questionnaire. This difficulty was primarily due to a lack of knowledge developed from the models used to train the system. A key aspect of the scenario used by Pracht was the fact that all of the "firms" were initially under-producing. This fact, once discovered and remedied, led to significant improvement in the health of the firm. The focus of this attention to details of producing and marketing the final product is reflected in these models and consequently in the knowledge base.

Table 4 System-generated versus Human Analysis of BML Variable Relationships

Variables		Sign	Number of subjects	
Market share	Sales volume	+	28	
Quality control	Market share	+	25	
Sales volume	Market share	+	28	
Back orders	Market share	N	21	-
Commission	Sales volume	- *	29	+
Production volume	Sales volume	0 *	22	+
Price	Net earnings	N	19	+ *
Quality control	Cost of goods sold	+	23	
Sales volume	Back orders	N	14	+ *
Sales volume	Production volume	+	20	
Sales volume	Raw material orders	+	25	
Sales volume	Net earnings	N	29	+ *
Advertising	Net earnings	+	15	
Production volume	Back orders	N	19	- *
Production volume	Raw material orders	+	26	
Cost of goods sold	Net earnings	N	27	- *
Plant capacity	Production volume	+	23	
Plant capacity	Cost of goods sold	+	19	

Notes: BML = Business Management Laboratory; “+” indicates belief that both variables change in the same direction; “-” indicates belief that the variables change in opposite directions; “0” indicates belief that there is no relationship at all; “N” denotes that SmartSLIM was unable to determine the relationship; * denotes the correct response where a discrepancy occurs.

Another potential explanation for this “gap” in the system’s knowledge is that the BML simulation was not run over an adequate amount of time for financial and administrative effects to force the subjects to incorporate this type of knowledge into their decision-making processes. Interestingly, the modal response for confidence in financial and administrative decisions by the student subjects in Pracht’s study was “neutral,” or a 3 on a scale from 1 to 5.

Table 4 shows the result of “administering” the second part of the questionnaire to the SmartSLIM system. In this part, the subjects indicated relationships between factors or elements of the BML game. A “+” indicates belief that both variables change in the same direction, a “-” indicates belief that the variables change in opposite directions, and a “0” indicates belief that there is no relationship at all. An “N” denotes that SmartSLIM was unable to determine the relationship.

In Table 4, the sign derived by the system appears in the column labeled “Sign.” The number and sign in the column labeled “Number of subjects” indicates the majority’s opinion in cases where the system did not reach the same conclusion. Where a discrepancy occurs, an asterisk denotes the correct response.

The system reached the same conclusion as (a majority of) the student subjects in over one-half of the cases that were capable of being tested. The system reached the correct conclusion (based on the BML code) in two-thirds of the cases tested. In one case, neither the system nor the students were correct.

Cases in which SmartSLIM and the subjects differed provide interesting insights into the methodology examined. In two of the seven cases of disagreement, the system was correct. While less than overwhelming, this indicates that in some cases the methodology may be at least useful for provoking a deeper reflection on the relationship in question when a disagreement occurs.

In three of the remaining five cases where there was disagreement, the relationship in question involved a financial entity, usually "net earnings." As mentioned earlier, financial relationships were found to be a weak area in the system's knowledge.

6. Summary

THE METHODOLOGY presented in Paradice and Courtney [8] has been implemented in the form of many new commands and procedures in the SmartSLIM decision support system. These enhancements can be categorized as facilitating (1) knowledge acquisition and knowledge base construction, (2) knowledge base perusal and manipulation, and (3) knowledge discovery. A taxonomy of relationships has been developed based on causal, non-causal, and quasi-causal categories.

An interactive method of knowledge acquisition has been implemented that captures a wide variety of information about hypothesized relationships. This information is stored in a knowledge base, the primary goal of which is to facilitate the analysis of "paths" from one variable to another. This analysis leads to an automatic synthesis of the information by the system, which leads to clearer understanding of the domain of variables on the part of the user.

The knowledge manipulation capabilities of the system also provide means of constructing models. The approach taken is one that allows previously untested models to be constructed by SmartSLIM. This provides another level of "discovery" by the system, since the system is not limited to testing only previously stored models.

A second feature of SmartSLIM's modeling capabilities is the ability to examine the knowledge base and act as an advisor. The system becomes an unbiased, expert consultant suggesting alternative actions for the decision maker to consider. Furthermore, the system allows the user to pursue alternative sub-goals in an interactive fashion as new relationships are discovered.

In conclusion, SmartSLIM has knowledge, at least to a limited degree, comparable to the student subjects. The system has also successfully determined some of the underlying relationships in the BML simulation. Although this implementation is not portable to other database systems, the methodology is independent of the characteristics of the SmartSLIM system. We conclude, therefore, that the methodology of acquiring, representing, and using knowledge investigated in this research is viable and that the methodology is a fruitful area for further research.

REFERENCES

1. Asher, Herbert B. *Causal Modeling*. Beverly Hills, CA: Sage Publications, 1983.
2. Blanning, Robert W. Management applications of expert systems. *Information and Management*, 7 (1984), 311-316.
3. Burns, J. R., and Winstead, W. H. An input/output approach to the structural analysis of digraphs. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-12, 1 (January/February 1982), 15-24.
4. Burns, J. R., and Winstead, W. H. Input and output redundancy. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-12, 6, (November/December 1982), 785-793.
5. Courtney, J. F., and Jensen, R. *SLIM User's Manual*. Dallas, TX: Business Publications, Inc., 1981.
6. Jensen, R. L., and Cherrington, M. L. *BML [Business Management Laboratory] Participant's Manual*. Dallas, TX: Business Publications, Inc., 1977.
7. O'Keefe, R. M. Expert systems and operational research—Mutual benefits. *Journal of the Operational Research Society*, 36, 2, 125-129.
8. Paradice, D. B., and Courtney, J. F., Jr. Controlling bias in user assertions in expert decision support systems for problem formulation. *Journal of Management Information Systems*, 3, 1, Summer 1986, 52-64.
9. Pracht, William E. GISMO: A visual problem-structuring and knowledge-organization tool. *IEEE Transactions on Systems, Man, and Cybernetics*, 16, 2 (1986), 265-270.
10. Pracht, W. E., and Courtney, J. F. A visual user interface for capturing mental models in model management systems. *Proceedings of the Nineteenth Annual Hawaii International Conference on Systems Sciences*, 1986, 535-541.