DYNAMIC CONSTRUCTION OF STATISTICAL MODELS IN MANAGERIAL DSS

David B. PARADICE and James F. COURTNEY

Department of Business Analysis and Research, College of Business Administration, Texas A&M University, College Station, Texas 77843-4217, USA

Abstract

The fields of operations research (OR) and artificial intelligence (AI) provide complementary methods that may be combined into managerial decision support systems (DSS). However, the management domain is substantially different from domains in which prior expert systems have been developed. Consequently, successful application of OR/AI techniques in managerial DSS requires careful analysis and additional development. Ongoing research concerning design and implementation of managerial DSS is discussed. A prototype system capable of constructing linear statistical models of direct and indirect relationships from a knowledge base of relationships is described and evaluated.

Keywords

Decision support systems, managerial expert systems, model construction, model management.

1. Introduction

The successes of expert systems in fields such as medicine, chemistry, and mathematics have generated speculation that similar results are obtainable in the field of management. In fact, success has been reported in applying artificial intelligence (AI) techniques in the areas of financial analysis (Bouwman [7,8]), portfolio management (Lee and Stohr [24]), general linear programming formulation (Murphy and Stohr [31]), managerial problem diagnosis (Courtney et al. [14]), and database design (Choobineh et al. [13]).

There are three distinct aspects of "knowledge engineering" which support the type of research mentioned above (O'Keefe [33]). These are: (1) the extraction of knowledge from human experts (knowledge acquisition), (2) the organization of knowledge into a knowledge base (knowledge representation), and (3) the use of the knowledge base in making reasoned deductions (knowledge usage).

According to O'Keefe, knowledge acquisition has been implicitly developed and practiced within operations research (OR) for some time. Further, knowledge representation and organization are modeling processes and thus familiar to the OR scientist, in approach at least ([33], p. 125). However, OR scientists may have little training in developing efficient and effective computer-based tools that begin to apply the knowledge in a knowledge base.

"Model management" has evolved as a research area which investigates methods of storing, modifying, and manipulating models. Model management research provides a natural arena for combining techniques from OR and artificial intelligence (AI). While the AI researcher can benefit from the experience of the OR scientist in model building, the OR scientist gains yet another tool for the OR "tool kit" [33]. Conceptual and theoretical work in model management can be found in Konsynski and Dolk [23], Elam and Henderson [17], Elam et al. [18], and Blanning [3–6].

Recent work by Lenard [25] has demonstrated a way of representing structured models in a relational database. This approach appears to hold great promise for model management because a wide range of model types may be represented as structured models. However, not all problems may be formulated in such well-defined formats. Moreover, determining the proper model structure may be a much more difficult task than solving the resulting formulation (McLean and Shepard [28]). Few would argue with the assertion that successful expert systems are predicated on an ability to determine and represent the proper model structure of the decision-making process at hand.

Indeed, the OR scientist works in a problem domain which is quite different from domains in which successful expert systems have been developed. The managerial problem domain is not "well-formed", as are the domains of physics, engineering, or mathematics – some of the areas in which expert systems have been developed. The managerial problem domain is composed of dynamic, temporal relationships between variables. Even the medical domain has the comparative advantage of time invariant relationships. Consequently, the design of a managerial DSS combining OR methods with AI techniques needs to be carefully considered.

The central issue of the research discussed in this paper is the design and implementation of managerial DSS in ill-formed rather than well-formed domains. This work focuses on the process of determining, constructing, and using proper model structure in ill-structured problem situations. In particular, we seek to develop a methodology which ultimately:

(1) is able to support formulation of ill-structured problems;
(2) is parsimonious with respect to knowledge representation issues;
(3) employs rigorous techniques from fields such as OR and statistics, where possible, to bring structure to ill-structured problem situations;
(4) incorporates a means to employ prior problem formulations to determine the structure of new problems;

(5) incorporates a way of explaining how a model structure is formulated;
(6) incorporates a means to accumulate knowledge about prior problem formulations in an advisory role, and
(7) is applicable to the general domain of management problems, as well as specialized domains such as finance and production.

The next section presents a brief review of the literature which emphasizes the importance of proper problem formulation. This is followed by a discussion of the attributes of the managerial problem domain that make construction of "expert systems" in management a particularly difficult task. We then present a taxonomy of relationship types that must be represented in a managerial "expert system". This is followed by a brief description of our prototype system in terms of the goals put forth above, and finally a few concluding remarks.

2. The problem of problem formulation

Although several models of the decision-making process exist, none is substantially different from Simon's [38] classic three-stage model consisting of intelligence, design, and choice. The intelligence stage is composed of environmental scanning and perception activity. The decision maker recognizes the need for a decision during the intelligence stage. The design stage consists of activity directed at creating alternatives that address the current decision-making situation. One alternative is selected in the choice stage.

The design and construction of most DSS, many of which employ MS/OR models, has been biased toward supporting the decision-making processes during the latter phases of the design stage and all of the choice stage. The ubiquitous DSS attribute, the WHAT IF question, certainly presupposes an existing model structure. A more serious problem is that it assumes a *correct* model structure. Although implementing AI techniques that will assist users in interpreting outputs of OR models will certainly yield benefits, a more crucial integration of the two fields may exist where AI techniques can be used to assist the user in formulating the correct OR model to be solved.

Ways of mitigating biases that may affect the construction of the correct model of a problem situation have been noted in earlier papers (Paradice and Courtney [34,35]). Any assumption that a correct model has been specified, especially if the model is new and does not have a history of producing satisfactory results, must be carefully examined. Consequently, we agree with Blanning [3] that problem diagnosis, or more generally, problem formulation, is an area where managerial expert systems could produce substantial benefits. An expert system approach to managerial problem diagnosis may be found in Courtney et al. [14].

3. Management's problem domain

Successful application of AI techniques has always been preceded by thorough knowledge of the (target) problem domain. This is frequently accomplished by extensive involvement of a "domain expert" in the development of the system. For example, an expert system in medical diagnosis is scrutinized during development by a highly skilled physician. An expert system in chemistry is similarly constructed under a chemist's supervision. Unfortunately, the characteristics of the domain in which management operates are quite different from domains in which many successful expert systems have been developed.

Traditionally, expert systems have been built for "deep and narrow" domains (Basden [2]). A deep and narrow domain is one in which the laws of the domain are well specified, such as the hard sciences, or some well-defined, man-made domains (e.g. tax laws). Management, on the other hand, is considered a "wide and shallow" domain. There is a wider range of potentially relevant factors in a wide and shallow domain, and there may be no specialists of the type found in deep and narrow domains.

There are several reasons why obtaining successful expert systems in wide and shallow domains may be difficult: distinguishing relevant from irrelevant factors may be difficult, testing the system exhaustively may be difficult, and achieving "real-time" response due to the size of the knowledge base required in such domains may also be difficult [2]. Moreover, the structure of the domain itself may change dynamically, resulting in continual updates to the knowledge base.

Basden points out ([2], p. 472) that such arguments may not be valid. First, to state categorically that such systems cannot be built is unwise, especially considering this is exactly what Japan has set out to do in its Fifth Generation computing research project (see Feigenbaum and McCorduck [20]). Second, in a wide and shallow domain, performance itself may not be very good; hence the system needs only to achieve better results than a human. Finally, there are situations where the need for completeness and accuracy are reduced. If the system can help users to better understand *aspects* of the problem domain, that may be justification enough for the construction of the system.

Research by Kasper [22], Pracht [37], and Loy [27] has demonstrated the benefits of modeling activity in understanding complex business problems. Modeling, we have noted, is at the foundation of much of the OR scientist's activity. Consequently, exploring ways of incorporating statistical techniques into managerial DSS to support the initial phases of problem formulation seems worthwhile.

4. Representation strategy

To begin the development of our general approach, we make the same observation as Lenard [25]: mathematical models are based on functional relationships.

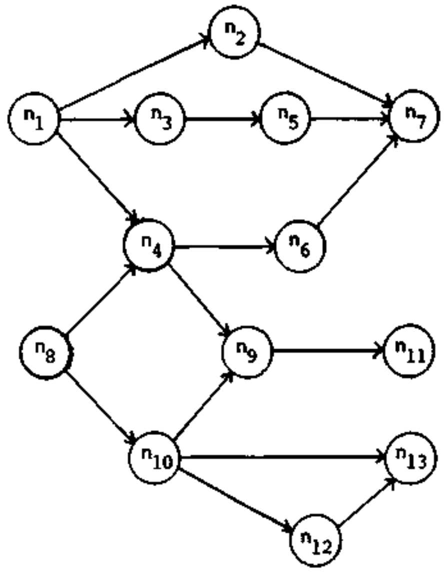

Fig. 1. An example of a structural model.

However, a representation strategy which is more suited to our needs is one closer to the realm of structural (not structured) modeling techniques. Hence, we begin with an expanded graphical representation based on the semantic network concept described by Elam et al. [18]. This representation is called a Structures Inheritance Network (SI-Net).

An SI-Net is composed of nodes and links for describing concepts and the interrelationships between these concepts. As described by Elam et al. ([18], p. 103), a "concept" is a set of functional roles tied together with an explicit structuring relationship. A concept is defined by two types of relationships. The first can be labeled "is a part of", which is represented by a definitional attribute link and describes the functional roles in a concept. The second type of relationship can be labeled "is structured as", which is represented by a structure link and describes how the roles are put together.

Figure 1 shows a conceptual view of one possible substructure in the problem domain (ignoring attributes). From this structure, we can deduce that n_7 is a function of direct relationships with n_2, n_5, and n_6, and that the variables n_1, n_3, n_4, and n_8 have an indirect relationship with n_7.

Following Newell [32], this approach can be characterized as very general, but weak. Although the existence of relationships between the variables can be deduced, little has been accomplished in terms of establishing the characteristics of the relationships. More information about the relationships is needed to increase the power of the model's ability to characterize the relationships.

For example, if causality and the usual regression assumptions are added as information, the relationships can then be characterized as the following system of equations:

$$n_7 = p_{27} n_2 + p_{57} n_5 + p_{67} n_6 + v_1 e$$

$$n_2 = p_{12} n_1 + v_2 e$$

$$n_5 = p_{35} n_3 + v_3 e$$

$$n_6 = p_{46} n_4 + v_4 e$$

$$n_3 = p_{13} n_3 + v_5 e$$

$$n_4 = p_{14} n_1 + p_{64} n_6 + v_6 e,$$

where $v_i e$ ($i = 1, \ldots, 6$) represent error components (Asher [1]). Hence, a much more specific (powerful) solution to the problem of characterizing the problem domain is obtained, but the generality under which the solution is applicable is sacrificed. If the assumptions that lead to the solution are reasonable, however, one is usually very willing to sacrifice generality.

A prototype system is described below which attempts to apply AI techniques in ways that translate a user's mental model of a problem domain from one in which only general (but weak) methods apply, to one in which powerful statistical tools may be applied. First, key aspects of the problem domain structure must be identified and maintained. This approach has an added advantage in that expert systems based on causal or "deep" knowledge have proven to be superior to those based on surface knowledge (Davis [15], Davis et al. [16], Chandrasekaran [11], Chandrasekaran and Mittal [12], Hart [21], Michie [29]).

5. A taxonomy of representation requirements

In order to operate effectively in a wide and shallow problem domain such as management, a managerial expert system must be able to represent the basic relationships between variables in the managerial problem domain. The first step in developing the ability to represent basic relationships is to categorize the relationships themselves. Three major taxonomic categories of basic relationships have been initially identified: causal, quasi-causal, and non-causal (Paradice and Courtney [36]). This taxonomy is based on experiences gained while developing a managerial expert system called SmartSLIM (described later) for a simulated management domain.

A relationship is causal when it is described by the following three conditions. First, the occurrence of one event must precede the occurrence of the second event. Second, there must be a relationship which can be formulated between the data for the two events. Third, there must be some justifiable theory or rationale for believing the relationship is causal. (See Einhorn and Hogarth [19] for a more detailed discussion of the nature of causal relationships.) For example, increased advertising expenditures are frequently believed to cause increased sales volume. Let $O_t(X)$ represent the observation of the variable X in time t, and $f(X)$ represent any function of the variable X. Then the causal relationship can be stated formally as

$$O_t(Y) = f(O_{t-1}(X)),$$

assuming a theoretical basis for the relationship.

Non-causal relationships (in this taxonomy) are strictly correlational. A relationship is correlational when data for the variables have a purely mathematical relationship, devoid of any cause or effect implications. In this situation, each variable takes on new values in a predictable fashion as the other variable changes value. If an organization tends to increase advertising expenditures each time it increases production volume, then these items are correlated.

A relationship is quasi-causal if a mathematical function describes the relationship between one variable and others in the problem domain. Three special cases of functional relationships may be defined in the business domain: bounding, definitional, and compositional.

A relationship is bounding when one variable acts as either an upper or lower bound on the value of a second variable. Plant capacity acts as an upper bound on production volume. Finished goods unit cost acts as a lower bound for price (assuming an ongoing concern) because price must be at least as great as the cost of production.

A definitional relationship is essentially 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, are variables that are categorized as definitional.

If the relationship is compositional and linear, and all of the compositional elements are known, then the following relationship holds for a variable X which is composed of n variables (denoted by y's):

$$X = \sum_{i=1}^{n} k_i y_i .$$

One must not be misled by the apparent simplicity of this type of relationship. Libby [26] and Courtney et al. [14] have demonstrated that linear models such as this are rather robust and widespread in management.

Figure 2 shows an adaption of the SI-Net concept representing a subset of the knowledge base of the SmartSLIM managerial problem formulation system. Attributes have been refined to distinguish graphically between subjective and calculated attributes. Structure links have been refined in terms of the taxonomy presented earlier.

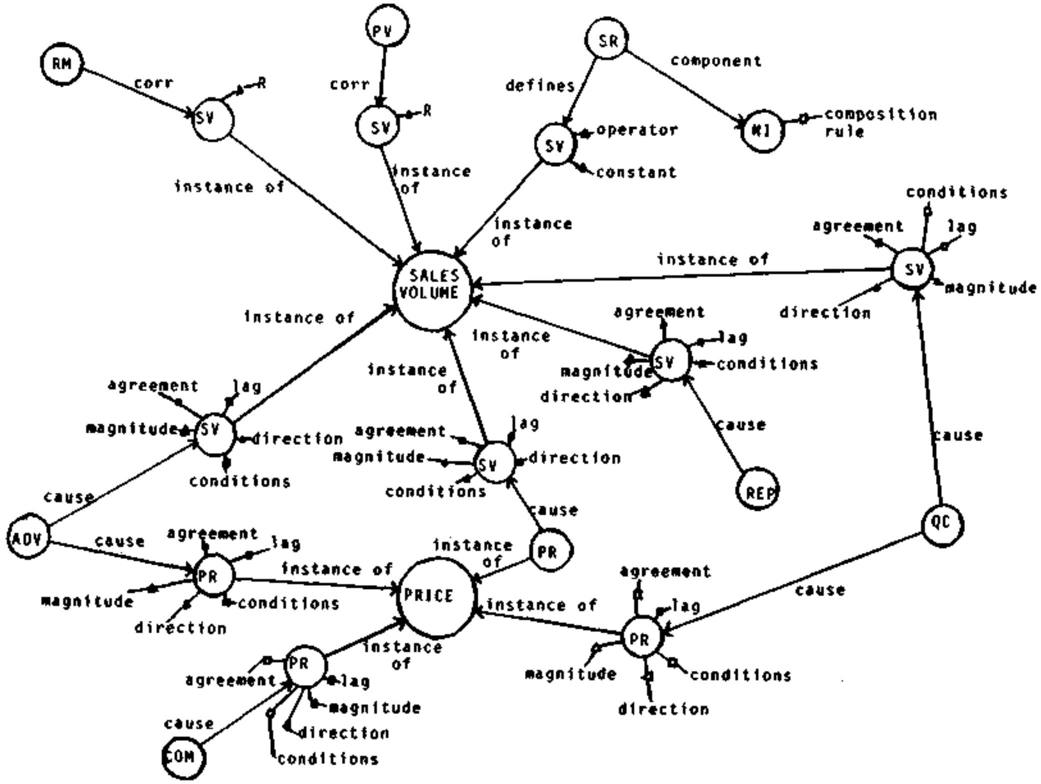

Legend:

ADV - advertising expense
COM - sales commission
NI - net income
PR - price
PV - production volume
QC - quality control
REP - sales representatives
RM - raw material inventory
SR - sales revenue
SV - sales volume

Fig. 2. A subset of the SmartSLIM knowledge base.

Structure links are: "causes", "is correlated with", "is a component of", and "defines". An additional structure type, "is an instance of", is maintained to support navigation through the SI-Net.

Each structure type has a template which specifies the data attributes required for full specification of the relationship. A correlation structure requires only the correlation (a calculated attribute) for full specification. A definition structure needs two calculated attributes: the numeric parameter for conversion and the operator required to perform the conversion. A composition structure has many attributes in the SmartSLIM system, including user-specified (i.e. subjective) attributes describing user agreement, time lag, and conditions under which the relationship holds, as well as calculated attributes involving the magnitude and direction (i.e. increasing or decreasing) of the relationship.

6. SmartSLIM: An expert system for managerial problem formulation

The preceding discussion is based in part upon strategies investigated during development of an expert system that could ultimately provide support for problem formulation activity in the management domain. The conceptual model for SmartSLIM is shown in fig. 3. In the discussion that follows, the focus is primarily on construction

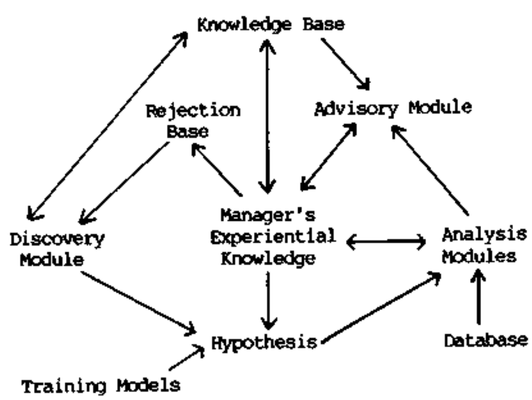

Fig. 3. Conceptual model of the SmartSLIM system.

of linear models by the SmartSLIM system. Although obviously a limited subclass of general OR models, this focus is sufficient to demonstrate the applicability of the methodology for acquiring, representing, and manipulating managerial problem domain modeling knowledge.

The SmartSLIM system "learns" through a series of sessions in which users (i.e. managers) hypothesize and test relationships between data items in an organizational database. For example, the CAUSE command may be used to input hypothesized

```
> CAUSE SV(F1-*),PRICE(F1-*),ADV(F1-*);

THE VARIABLES RETAINED IN THE MODEL ACCOUNT FOR 73.5% OF THE VARIATION IN SV.

A ONE UNIT INCREASE IN PRICE RESULTS IN A -8573.3 DECREASE IN SV.  A MORE
EXTREME COEFFICIENT FOR PRICE WOULD BE OBTAINED ONLY 3.58 TIMES IN 100,000,000.

SHOULD THIS RELATIONSHIP BE STORED FOR FUTURE REFERENCE? Y

ON A SCALE FROM 1 TO 10 WHERE 10 INDICATES STRONG AGREEMENT, HOW DO YOU AGREE
WITH THIS? 8

HOW MANY QUARTERS PASS BEFORE PRICE AFFECTS SV? 0

IS SV LIMITED BY THE SIZE OF PRICE? N

DOES A MINIMUM VALUE OF PRICE GUARANTEE SOME MINIMUM LEVEL OF SV? N

IS THIS RELATIONSHIP DEPENDENT ON OTHER CONDITIONS? N

ENTER A DESCRIPTION OF THIS RELATIONSHIP:
INCREASES IN PRICE RESULT IN DECREASES IN SALES VOLUME
```

Fig. 4. An excerpt from the dialogue examining the hypothesis that price and advertising impact (cause) sales volume.

causal relationships between one dependent variable and one or more independent variables. An example of dialogue involving the CAUSE command is shown in fig. 4. The CAUSE command determines the best linear regression model of these variables using Mallows' Cp criterion (Montgomery and Peck [30]). Mallows' Cp may be interpreted in terms of parameter precision and bias. When the parameter estimates are unbiased, the value of Cp will be nearly equal to the number of variables in the model. Hence, low values of Cp that are close to the number of variables in the model indicate parsimonious models with small bias. After statistical analysis of the hypothesized relationship, the user has the opportunity to express a level of agreement with the system's calculations. Over time, a knowledge base is created (as shown earlier in fig. 2) which contains information about relationships in the managers' domain. Later, any manager can query the system in order to obtain guidance regarding a problem being faced. Thus, the system may be used to communicate "expert" knowledge to novice managers or trainees.

This manager–SmartSLIM interaction can now be illustrated. Typically in the use of computer-based systems, management personnel validate a system by checking that certain obvious outputs are generated. For example, sample cases with known results may be input to a simulation package to verify that the package produces the correct result. In the case of a managerial system, the system would be required to demonstrate knowledge of certain fundamental relationships. This validation process is illustrated in fig. 5, where the user checks the system's knowledge of the relationship between price and sales volume. SmartSLIM uses causal modeling with path analysis (Asher [1]) to synthesize the knowledge in the SI-Net. Path analysis attempts to reconstruct the correlation between variables by examining the weights of the paths between the relationships. Path analysis is initiated by the PATH

```
> PATH PRICE, SV;
WOULD YOU LIKE EXPLANATIONS (Y/N)? N

PRICE -> SV
CONTRIBUTION: -0.44682
NOTE: THIS DIRECT PATH CONTRIBUTION IS NOT INCLUDED
      IN THE TOTAL.

TOTAL CONTRIBUTION OF INDIRECT PATHS IS 0.00000
THE CORRELATION BETWEEN PRICE AND SV IS -0.44682
NO INDIRECT PATHS.

THE DIRECT PATH EXISTS IN THE KNOWLEDGE BASE
(AGREEMENT = 8).
```

Fig. 5. Stored knowledge of the relationship between price and sales volume.

```
> PATH COGS, PRICE;
WOULD YOU LIKE EXPLANATIONS (Y/N)? N

COGS -> RMO -> RMI -> PVFS -> PRICE
CONTRIBUTION: -0.26881

COGS -> RMO -> RMI -> FGUC -> PRICE
CONTRIBUTION:  0.02573

COGS -> RMO -> RMI -> PVFS -> FGI -> PRICE
CONTRIBUTION:  0.06827

TOTAL CONTRIBUTION OF INDIRECT PATHS IS -0.174799
THE CORRELATION BETWEEN COGS AND PRICE IS -0.45539
E - I ANGLE = 26.56505 DEGREES.
ON A SCALE OF 1 TO 10 WHERE 10 IS STRONG BELIEF,
THE SYSTEM ASSIGNS A VALUE OF 3 TO THE OVERALL
CONCLUSION THAT INCREASES IN COGS INCREASE PRICE.
```

Fig. 6. Constructed knowledge of the relationship between cost of goods sold and price.

command. Here, SmartSLIM has indicated that price and sales volume are believed to be negatively related. Other fundamental relationships would be validated similarly by users.

Once convinced that the system has some basic knowledge of the business domain, a user might then instruct the system to construct more complex statistical models. The user may desire to determine what the system knows regarding the relationship between cost of goods sold and price. This could be initiated by another execution of the PATH command, as in fig. 6.

Here, SmartSLIM has no direct knowledge of the relationship of cost of goods sold to price. The system has, however, constructed a relationship through the indirect paths between these two variables. Not only has SmartSLIM been able to construct a relationship, it also gives an indication of the direction of the relationship by analyzing the signs of the sequences between the variables (Burns and Winstead [9,10]). Interestingly, in this case the system has determined that the direction of the relationship is opposite to the direction indicated by the correlation. The discrepancy between the amount of the path contribution and the correlation is an indication

```
> 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
```

Fig. 7. Advisory action of system regarding goal of increasing sales volume.

that the knowledge base needs more training in this area. This need is also reflected in the low confidence level assigned by the system.

The confidence level indicates SmartSLIM's level of confidence 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 most suspect. SmartSLIM takes the strongest of the weakest agreement levels (i.e. a maximin approach) to be its confidence in its calculations.

Finally, a user may wish to have the system construct a statistical model incorporating all of the system knowledge regarding sales volume. This is done in fig. 7, using the MODEL command. In this case, the user indicates an interest in increasing sales volume by specifying the optional INCREASE keyword.

Here, SmartSLIM has once again gone to the knowledge base and dynamically constructed a model. First, variables that are related to sales volume in a non-causal manner are reported. Then, any variable known to directly influence sales volume has been included in a model. A stepwise regression approach is taken again to determine the best model. SmartSLIM presents its estimate of the coefficients in the model, giving the user a current (as of the most recent database update) estimation of the relationships involved. Subsequently, any polynomial relationships are constructed and presented.

Since the user specified the INCREASE keyword, the system has generated a list of variables that potentially could be manipulated in order to achieve the goal of increasing sales volume. This list includes advertising expense, price, raw material inventory, and back orders among other variables.

The user indicates that more information about the relationship between sales volume and price is desired. SmartSLIM responds with detailed information about this relationship, which is drawn from the knowledge base. This information leads the user to examine ways of achieving a new goal: that of decreasing price. This is done by requesting more information about price. SmartSLIM now responds that variables that can be manipulated to achieve this goal include sales commission, advertising expense, and quality control.

In this case, the system has assumed an advisory role, assisting the user in learning about relationships in the business domain that are under user control. This type of active participation on the part of the system may prove to be one of its greatest assets.

7. Discussion

In order to place the work-to-date in perspective, the SmartSLIM system can be examined relative to the goals put forward in the introduction. Preliminary tests indicate that the approach employed in SmartSLIM is capable of supporting problem formulation in ill-structured problem domains. SmartSLIM's performance has been shown comparable to students' in deducing relationships in a simulated managerial problem domain (Paradice and Courtney [34]). In fact, SmartSLIM was successful at deducing some relationships that the majority of students mis-specified.

The knowledge representation scheme in SmartSLIM is parsimonious. Initial versions of the system attempted to limit the knowledge base to causal relationships.

This constraint resulted in unsatisfactory performance, leading to the development of the taxonomy of relationships presented above. The resulting taxonomic structure is still quite simple, but it has proved to be quite robust, also.

Methods which will support construction of more complex OR and statistical models than those represented here are being investigated. Key elements (time lags, constraining variables, etc.) needed for such model construction are captured during the knowledge acquisition sessions described earlier. The emphasis on linear models to date has been driven by concentration on resolving methodological issues related to knowledge acquisition, knowledge representation, and knowledge manipulation in managerial problem domains. Additionally, linear models are widely applicable in many problem domains (see Libby [26]).

SmartSLIM has an explanation facility which describes the relationships being manipulated as the (linear) models are being constructed. The system draws on prior problem formulations to solve new problems, and allows users to make ad hoc requests for advisory information. Although the SmartSLIM system interfaces with a database of approximately 125 organizational variables maintained for over thirty firms, the system's response time is very fast. Typical response times are 1 to 5 seconds.

Finally, the SmartSLIM system is applicable to the general domain of managerial problems. The knowledge representation structure has been developed to maintain the general types of relationships specified in the taxonomy presented above. This taxonomy is neutral with respect to functional areas in management.

8. Summary

The current version of the SmartSLIM system begins to bridge the gap between the expertise of OR scientists (knowledge acquisition and organization) and the expertise of AI researchers (automated knowledge usage). Unlike some other experts, the OR scientist works in a problem domain that is extremely ill-structured. As such, new approaches in design and implementation for managerial DSS must be developed in order to incorporate the methods of the OR scientist.

Future versions of SmartSLIM will be aimed at creating more complicated OR models dynamically. At least two key areas of research are anticipated to achieve this goal. First, the SI-Net must be augmented to model simultaneous systems of equations in a more direct manner. Secondly, a separate knowledge base of model types (and distinctions) will be needed.

References

[1] H.B. Asher, *Causal Modeling* (Sage, Beverly Hills, CA, 1983).
[2] A. Basden, On the application of expert systems, Int. J. Man-Machine Studies 19, 5(1983) 461.

[3] R.W. Blanning, Management applications of expert systems, Information and Management 7, 6(1984)311.

[4] R.W. Blanning, Issues in the design of relational model management systems, *Proc. National Computer Conference* (1983) pp. 395 – 401.

[5] R.W. Blanning, A relational framework for model management in decision support systems, DSS-82 Trans. (1982)16.

[6] R.W. Blanning, Information systems for model-based planning, *Proc. 10th Annual Conference – Society for Management Information Systems* (1978) pp. 132 – 140.

[7] M.J. Bouwman, Human diagnostic reasoning by computer: An illustration from financial analysis, Manag. Sci. 29, 6(1983)653.

[8] M.J. Bouwman, Towards expert systems: The analysis of expert financial behavior, paper presented at the ORSA/TIMS National Conference, Atlanta, Georgia (1985).

[9] J.R. Burns and W.H. Winstead, An input/output approach to the structural analysis of digraphs, IEEE Trans. on Systems, Man, and Cybernetics SMC-12, 1(1982)15.

[10] J.R. Burns and W.H. Winstead, Input and output redundancy, IEEE Trans. on Systems, Man, and Cybernetics SMC-12, 6(1982)785.

[11] B. Chandrasekaran, Towards a taxonomy of problem solving types, AI Magazine (Winter/ Spring 1983)9.

[12] B. Chandrasekaran and S. Mittal, Deep versus compiled knowledge approaches to diagnostic problem solving, Int. J. Man-Machine Studies 19(1983)425.

[13] J. Choobineh, M.V. Mannino, J. Nunamaker and B. Konsynski, The architecture of a rule-based system for database design based on forms, *Proc. 19th Annual Hawaii Int. Conf. on System Sciences* (1986) pp. 581 – 588.

[14] J.F. Courtney, D. Paradice and N. Ata, A knowledge-based DSS for managerial problem recognition and diagnosis, special issue of Decision Sciences on DSS 18, 3(1987)373.

[15] R. Davis, Reasoning from first principles in electronic troubleshooting, Int. J. Man-Machine Studies 19(1983)403.

[16] R. Davis, H.E. Shrobe, W. Hamscher, K. Wickert, M. Shirley and S. Polit, Diagnosis based on description of structure and function, *Proc. American Association for Artificial Intelligence* (1982) pp. 137 – 142.

[17] J. Elam and J.C. Henderson, Knowledge engineering concepts for decision support system design and implementation, Information and Management 6, 2(1983)109.

[18] J. Elam, J.C. Henderson and L.W. Miller, Model management systems: An approach to decision support in complex organizations, *1st Int. Conf. on Information Systems*, Philadelphia, PA (1980) pp. 98 – 110.

[19] H. Einhorn and R. Hogarth, Prediction, diagnosis, and causal thinking in forecasting, Journal of Forecasting 1(1982)23.

[20] E.A. Feigenbaum and P. McCorduck, *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World* (Addison-Wesley, Reading, MA, 1983).

[21] P.E. Hart, Direction of AI in the eighties, SIGART Newsletter 79(1982).

[22] G. Kasper and R.P. Cerveny, A laboratory study of user characteristics and decision-making performance in end-user computing, Information and Management 9(1985)87.

[23] B.R. Konsynski and D. Dolk, Knowledge abstraction in model management, DSS-82 Trans. (1982)187.

[24] J.B. Lee and E.A. Stohr, Representing knowledge for portfolio management decision making, Working paper, New York University (1985).

[25] M.L. Lenard, Representing models as data, *Proc. 19th Annual Hawaii Int. Conf. on System Sciences* (1986) pp. 389 – 396.

[26] R. Libby, *Accounting and Human Information Processing: Theory and Applications* (Prentice Hall, Englewood Cliffs, NJ, 1983).

[27] S.L. Loy, An experimental investigation of a graphical problem-structuring aid and nominal group technique for group decision support systems, Ph.D. Dissertation, Texas Tech University (1986), unpublished.

[28] J.M. McLean and P. Shepherd, The importance of model structure, Futures (1976)40.

[29] D. Michie, High road and low road programs, AI Magazine 3, 1(1982)21.

[30] D.C. Montgomery and E.A. Peck, *Introduction to Linear Regression Analysis* (Wiley, New York, NY, 1982).

[31] F.H. Murphy and E.A. Stohr, An intelligent system for formulating linear programs, Working Paper Series No. 95, Center for Research on Information Systems, New York University (1985).

[32] A. Newell, Heuristic programming: Ill-structured problems, in: *Progress in Operations Research*, Vol. III, ed. J.S. Aronofsky (Wiley, New York, NY, 1969).

[33] R.M. O'Keefe, Expert systems and operational research – mutual benefits, J. Oper. Res. Soc. 36, 2(1985)125.

[34] D. Paradice and J.F. Courtney, Causal and non-causal relationships and dynamic model construction in a managerial advisory system, Journal of Management Information Systems 3, 4(1987)39.

[35] D. Paradice and J.F. Courtney, Controlling bias in user assertions in expert decision support systems for problem formulation, Journal of Management Information Systems 3, 1(1986) 52.

[36] D. Paradice and J.F. Courtney, SmartSLIM: A DSS for controlling biases during problem formulation, in: *Human Factors in MIS*, ed. J. Carey (ABLEX, Norwood, 1987), forthcoming.

[37] W.E. Pracht, A graphical interactive structural modeling aid for decision support systems, IEEE Trans. on Systems, Man, and Cybernetics 14(1986).

[38] H.A. Simon, *The New Science of Management Decision* (Harper and Row, New York, NY, 1960).