Expert System for Aspects of the TSM Process

M. A. MORRIS AND L. J. POTGIETER

An expert system technology is applied to the transportation system management (TSM) process. Two simple expert systems were built to test the feasibility of applying expert system technology to certain aspects of TSM. Findings of the feasibility study were then used to define requirement specifications for a practical expert system to provide multisolutions for problems identified through the TSM process. Steps taken to implement a full-scale TSM expert system are then described. Experience gained during the feasibility stage of the project indicates that successful knowledge engineering is the key to successful expert system development. An extensive and detailed description of the formal knowledge engineering approach used to clarify and structure the TSM knowledge into a form usable in an expert system is provided. To test the effectiveness of the expert system, output from the system was compared with recommendations made by transportation consultants on seven large-scale intersection problems. The test indicates that the output of the system compares favorably with the recommendations made by human experts.

The systems approach to problem solving in transportation is well established and has been effective in optimizing existing transportation infrastructures and operations. However, limited funds and expertise have placed constraints on this method and it has become apparent that new tools and formalized methodologies are required to make more effective use of this approach.

An expert system is constructed for certain aspects of the transportation system management (TSM) methodology. The expert system is intended to be a practical working system to be used when necessitated by a lack of TSM expertise.

Practical transportation engineering knowledge is acquired and structured into a form usable by an expert system shell (called knowledge engineering by expert system builders).

Knowledge engineering is becoming the underlying factor in the success or failure of expert system development (*1*). Chang (*2*,*3*), Maher (*4*), and others have discussed the purely technical details of expert system construction relating to transportation.

BACKGROUND

The TSM approach has been adapted from the extensive TSM literature to suit local conditions and constraints (*5–7*).

Tomecki (*6*) defines the TSM process as a seven-stage process as follows:

- Stage 1. Public communications of improvement needs and potential,
- Stage 2. Problem definitions,
- Stages 3 and 4. Generation and analysis of alternative solutions,
- Stage 5. Evaluation and selection of preferred alternative solutions, and
- Stages 6 and 7. Implementation and monitoring.

Stages 1 and 2 are well defined and understood (*7*). A step-by-step procedure has been developed to obtain, through public participation, a range of problem definitions.

Stages 3 and 4 are less straightforward. Extraneous factors, such as the shortage of TSM expertise and the fragmentation of TSM techniques throughout the literature, have affected the generation of effective and wide-ranging solutions to transportation problems brought to light through Stages 1 and 2 of the TSM process. In order to overcome these problems, an expert system was proposed to address Stages 3 and 4 directly.

For each problem identified during Stages 1 and 2 of the TSM process, the expert system would be required to generate a range of solutions (Stage 3). The expert system would assist with the initial analysis (Stage 4) of the generated solutions by providing a weighted certainty factor (*8*) as to the likely effectiveness of such a solution.

These solutions are seen as proposals for further investigation and analysis and serve as a guide to multiple solutions for a given problem.

TSM EXPERT SYSTEM REQUIREMENT DEFINITION

TSM Expert System Feasibility Study

Two simple expert systems were built to investigate the feasibility of implementing an expert system for TSM. The objectives of the feasibility study were to investigate

- The method of representing transportation knowledge (the knowledge representation) that would be most suitable for the TSM project,
- The type of user interface that would be most appropriate, and
- The applicability of expert system technology to transportation planning.

Expert system shells were used for the two systems, including the KES II production system (*9*) (using a production rule or If . . . Then . . . form of knowledge representation) and

M. A. Morris, Transportation Research and Consultancy, P.O. Box 95230, Waterkloof, 0145, Pretoria, South Africa, L. J. Potgieter, Town Council of Springs, P.O. Box 45, Springs, 1560, South Africa.

the KES II hypothesis and test (*10*) (using a frame-like form of knowledge representation). These two shells were chosen because they provided different forms of knowledge representation, but were similar in all other respects.

Transportation information needed for both systems was obtained from the tables within the *Simplified Procedures for Evaluating Low-Cost TSM Projects* (*5*).

Construction of these feasibility study expert systems took place over a 4-month period. When completed, both expert systems were able to perform at an acceptable level and a decision was made to implement the TSM expert system.

TSM Expert System Requirement Specifications

The feasibility study gave valuable practical insight into applying expert system technology to the TSM process. The following requirements were specified for the TSM expert system on the basis of the experience gained during the feasibility study:

1. The production rule form of knowledge representation would be used for the TSM expert system. In practice, it was found that both of the methods for representing transportation knowledge discussed previously were effective in modeling transportation problems. However, the production rule method was found to be more understandable by the transportation engineers involved in the project.
2. Type and form of input to the expert system should be clearly specified. For example, information, such as traffic flow, may be given in numeric terms as vehicles per hour or as symbolic values such as low, medium, high, or saturated. Availability of such information for a given problem area or site also needs to be taken into account.
3. Type and form of the output should be clearly specified and all recommendations made by the system should be clearly understood. The feasibility study systems indicated that recommendations made by the expert systems were found to be ambiguous and were not well understood by users.
4. An intelligent front end program should be added to the expert system. The feasibility study found that lengthy question-and-answer sessions between the expert system and a user (typically 30 to 40 questions) often resulted in confusion on the part of the user. This problem was overcome by adding a program to the expert system to assist with the initial capture of information.
5. The sequence of questions asked by the expert system should, as closely as possible, mimic the question sequences and style of a typical transportation engineer or human expert. The feasibility study system contained no instructions for controlling the sequence of questions asked by the expert system. Questions tended to be presented to the user in an illogical sequence. Subjective performance of the expert system improved dramatically when structures (available in many expert system shells) were used to order the question sequence.

In addition to these requirements, the feasibility study demonstrated the importance of a clearly defined knowledge engineering methodology.

STRUCTURING AND ANALYZING THE TSM KNOWLEDGE

The following discussion focuses on the structuring and analysis of TSM knowledge and the development of production rules (If . . . Then . . . statements) for use in the expert system shell. This process, called knowledge engineering, is critical to the success of an expert system project. It is essential that a coherent approach is used to obtain and structure the knowledge within the domain to be modeled (*10,11*).

Modeling the Transportation Engineer's Approach to Problem Solving

During the feasibility study, an intense 1-day session was held to identify how transportation engineers use the TSM process to assist in the solution of transportation problems. The session also served to familiarize the expert system builder with transportation concepts.

The discussion showed that transportation engineers often use a broad two-step process when using TSM for problem solving.

Step 1. Overview Questions

Overview questions were used to obtain general information on the type and location of the problem. First, engineers required information on the location of the problem. This information was used to choose an appropriate problem category (e.g., isolated intersection, corridor, and employment center). Next, within each problem category, information was required on specific topics relevant to the problem area. For example, when the problem category was isolated intersection, the engineer would require information on geometric layout, traffic condition, problem symptom, etc. Discussions indicated that there was a specific set of topics for each problem category.

An example of the type of questions asked during the overview question step would be: "What form of traffic control is being used at the intersection?"

Step 2. Detailed Questions

Information obtained during the overview questioning was then used to guide the engineer in asking specific detailed questions relevant to the problem under consideration. An example of the type of questions asked during the detailed question step would be: "You have said that queues are forming in the left-hand lane at the intersection. Does this happen throughout the day or only at peak times?"

Structuring TSM Knowledge

In order to provide structure for the body of knowledge to be modeled in an expert system, it is necessary for the expert system builders to make a thorough study of the broad areas making up the knowledge domain. Much of the early work is devoted to this process. Fortunately, in TSM methodology much of the work on the structuring of the knowledge domain was directly available in the TSM documentation. TSM documentation (*5*) proved to be an excellent source for much of the knowledge engineering work of the project.

In order to mimic the approach followed by engineers, the knowledge area or domain was divided into seven problem categories:

- Isolated intersections,
- Street segments,
- Corridors,
- Residential communities,
- Employment centers,
- Commercial centers, and
- Regional, system-wide.

Each problem category was then treated as a separate expert system. To date, isolated intersection and street segment systems have been completed. Work is continuing on the other five expert systems.

Once work began on isolated intersections, discussions were held with transportation engineers to identify the major topics or areas of required information (ARIs) for which overview information was required. The following ARIs were identified:

- Geometric layout,
- Traffic control,
- Traffic conditions,
- Traffic problem symptoms,
- Pedestrian conditions,
- Pedestrian problem symptoms, and
- Actions.

In order to ensure that the system followed, wherever possible, the natural question sequence used by a human expert, an informal dependency graph was developed for the ARIs within each problem category.

The dependency graph idea has proved to be useful in diagrammatically representing the dependency relationships between the various ARIs within a problem category. Dependency graphs were used to describe the usual question sequence used for obtaining ARI information. Figure 1 shows the complete dependency graph for the problem category of isolated intersection. Within this dependency graph, it can be seen that intersection control is dependent on intersection traffic condition and intersection geometric layout.

Analyzing the TSM Knowledge

After all the ARIs for a particular problem category were defined, the knowledge required for each ARI was analyzed and placed in a form acceptable to the expert system shell used for the project. All possible values for each ARI were then listed. A similar approach is outlined in Weiss and Kulikowski (*12*). Formally, the ARI values may be seen as elements of each ARI set. For intersection geometric layout, the ARI values are

- Slip lane,
- Left-turn storage bay,
- Median—open,
- Median—closed,
- High-occupancy-vehicle lanes, and
- Unorthodox layout.

ARI values for intersection control are

- No control,
- Stop signs,
- Yield signs,
- Traffic lights—fixed time,
- Traffic lights—semiactuated, and
- Traffic lights—fully activated.

Because the actions for each problem category may also be seen as an ARI, the ARI values (list of all possible actions for a problem category) were compiled from the TSM documentation (and augmented by actions relevant to local conditions). These ARI value lists formed the basis for the production rules (If . . . Then . . . statements) used in the knowledge base of the expert system.

Examples of actions taken from the ARI value list for the problem category of isolated intersection are as follows:

- Add a left-turn storage bay,
- Add a right-turn lane,
- Add a left-turn arrow phase, and
- Upgrade intersection layout.

The complete list of actions that were considered was compiled from the broad TSM literature.

The analysis to be discussed and the structuring discussed previously were then used to develop the production rules or If . . . Then . . . statements used by the expert system. Figure 2 shows an example production rule developed from the

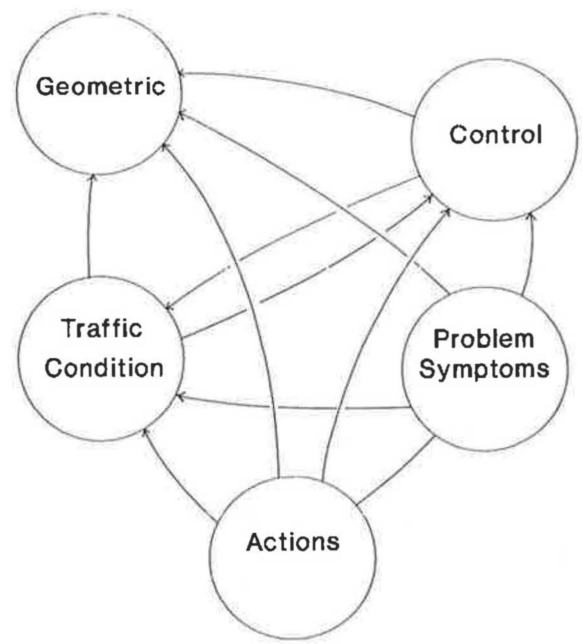

FIGURE 1 ARI dependency graph for ISOLATED INTERSECTION.

isolated intersection problem category that illustrates the concepts previously discussed. In this example, the isolated intersection ARIs, ordered per the dependency graph, were geometric, traffic control, traffic condition, traffic symptom, pedestrian condition, and pedestrian symptoms. For each of these ARIs, ARI values (i.e., no left-turn storage bay, any, saturated or high or medium, etc.) were used in developing the production rule example. The outcome action was obtained from the action list for isolated intersections.

Once the basic rule structure was defined, it was possible to systematically acquire the knowledge relating to the area being modeled. The following discussion describes the steps used to develop the detailed If . . . Then . . . information.

Eliciting the Detailed Rules

Interviews were held with experienced transportation engineers to define the rules in the form given previously. The production rules were handled one at a time. Two engineers were involved in each interview session. A third member of the expert system building team acted as a facilitator. Engineers were presented with an empty rule as shown in Figure 3. The engineers were then asked to complete the rule. The ARI value list for each ARI was used as a guideline for completing the rule.

A flexible approach was adopted. Where no value seemed appropriate to the rule being considered, new information was added to the ARI value lists. In several cases during the interview sessions, it was found that the general ARI information was inadequate for selecting a specific action. Once the ARI information was in place, the discussion group identified any specific information that they felt was specific to the rule under discussion. Specific information for the example would be whether space is available for adding a left-turn lane. A possible question generated by the expert system would be: "There are indications that a left-turn storage bay would improve the problem at the intersection. Is there room to add a left-turn lane?"

1. YES,
2. NO.

Knowledge Structuring and Analysis Summary

The underlying operation of an expert system is beyond the scope of this discussion. However, each element in a production rule is closely linked to a question generated by the expert system. For example, geometric layout would generate a question such as: "Which of the following describe the geometry of the intersection?"

1. Slip lane,
2. Left-turn storage bay,
3. Median,
4. HOV lanes, and
5. Unorthodox layout.

The expert system user would respond appropriately.

The rule structure is related to modeling of the transportation engineers' approach to problem solving as follows:

- ARIs in each rule are used to generate overview questions,
- Specialized information is used to generate detailed questions for the specific rule, and
- ARI order is related to the order obtained from the dependency graph and controls the question sequence, which is expert system shell-specific.

```
IF
 GEOMETRIC          = NO  Left turn storage bay      AND
 TRAFFIC CONTROL    = Any                            AND
 TRAFFIC CONDITION  = Saturated or High or Medium    AND
 TRAFFIC SYMPTOM    = Queues in left turn lane
                      impeding straight traffic      AND
 PED CONDITION      = Any                            AND
 PED SYMPTOMS       = Any                            AND
 SPECIFIC INFORMATION
THEN
 ACTION = Add a left turn storage bay.
```

FIGURE 2 Production rule developed from ISOLATED INTERSECTION—with ARI values.

```
IF                                                  F3
 GEOMETRIC          =                                AND
 TRAFFIC CONTROL    =                                AND
 TRAFFIC CONDITION  =                                AND
 TRAFFIC SYMPTOM    =                                AND
 PED CONDITION      =                                AND
 PED SYMPTOMS       =                                AND
 SPECIFIC INFORMATION =
THEN
 ACTION = Add a left turn storage bav.
```

FIGURE 3 Production rule developed from ISOLATED INTERSECTION—without ARI values.

OPERATION OF THE TSM EXPERT SYSTEM

The TSM expert system program has been designed to run on an IBM AT PC or compatible computer. At start up, the user is presented with the list of standard TSM problem categories and the user is required to choose an appropriate category.

The system responds with a form and diagram on the screen. The structure of both the form and the diagram depends on the problem category chosen. In the case of isolated intersections, for example, the user is presented with a stylized diagram of an intersection as shown in Figure 4.

Each solid rectangle indicates a location where a pop-up window can be activated to enter relevant information (for example, traffic flow and pedestrian flow) into the diagram. Geometric, traffic, and pedestrian flow information is entered into the diagram (through a series of pop-up menus). Information is entered separately for each approach to the intersection. A similar diagram form is used to enter problem symptom information for each approach to the intersection.

On completion of the diagram and form, the expert system proceeds with a series of specific questions (the detailed questions). Once adequate information has been obtained from a user, the system responds with a series of recommended actions, which are given in a provisional order.

The complete expert system consists of three interacting subsystems:

1. The KES II production system expert system shell,
2. The Turbo Prolog intelligent front end, and
3. Supporting C language functions.

The KES II system is embedded within the C master program. The Prolog program has been compiled separately, but is executed via a system command from the (master) program.

At startup, the user is given an option to run the expert system using data from

1. A previously stored and named case study, or
2. The previous consultation.

Alternatively, a new consultation can be initiated. The user is then asked to choose the broad problem category. The master C program then runs the Prolog program, which displays the appropriate diagram on the screen.

Because the diagram and ancillary questions such as geometric layout information have been implemented on a virtual screen, it is possible for the user to move around the screen via the edit keys. The virtual screen contains diagrams and entry fields for data required by the system. The F10 key terminates this phase of the program. A second similar diagram, for problem symptoms, is then displayed on the screen. Once again, F10 terminates the input phase of the program.

Information gathered by the Prolog program is then written to a file, which is automatically read by the KES expert system shell.

The program then exhibits typical expert system behavior. Users are asked a series of questions on the basis of the initial input from the diagrams. Conclusions are then displayed on the screen. Users can opt to repeat the consultation, terminate the consultation, or store the data from the entire consultation as a named case study.

KES II PS

The KES II PS shell, using the standard backward chaining inference strategy, was used as the basis of the TSM expert system. The system provided several features that were found to be useful.

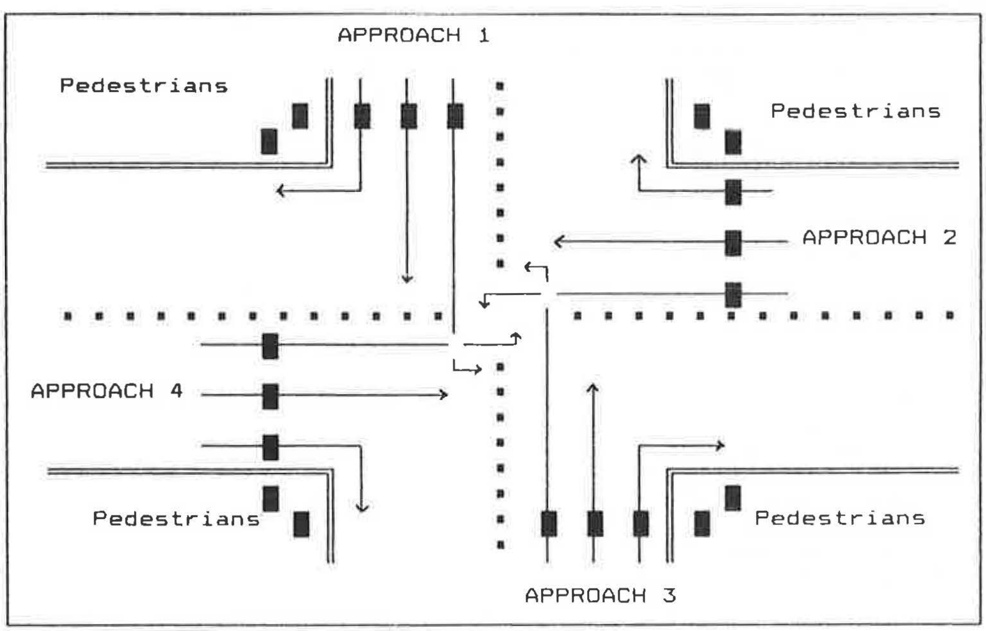

FIGURE 4 TSM expert system stylized diagram of an intersection for ISOLATED INTERSECTION.

• The KES system could be embedded in the C program and viewed as a series of functions to be called when necessary,

• Multiple knowledge bases could be loaded and unloaded from within the C program, and

• The shell has simple, yet powerful file-handling capabilities for transparently interacting with other programs within the DOS environment.

Turbo Prolog

Turbo Prolog 2.0 and the Turbo Prolog Toolbox were used to develop the virtual screens and windows used by the expert system. The Prolog language with its pattern matching ability provided an excellent basis for the development of intelligent forms and diagrams. In addition, the Prolog program was designed to prevent the user from entering contradictory information into the system.

C

Lattice C V 3.01 formed the basis for the complete expert system. It was chosen to maintain compatibility with the C interface to the KES system.

VALIDATING THE EXPERT SYSTEM

Ongoing Validation

Validation and testing of expert system performance is an integral part of expert system construction. The knowledge within the TSM expert system is based on a human expert's interpretation of a given situation or problem and therefore cannot be assumed to be 100 percent correct (*13*).

Physical construction of the TSM expert system was an iterative process—a small number of production rules was added to the system and then adjusted and tested until the system produced satisfactory results. Procedures were repeated until all the rules were added to the system.

However, it was felt that this approach did not provide exhaustive testing of the expert system.

Practical Validation

In order to ensure that the expert system provided useful practical results, the expert system's recommendations were compared with recommendations obtained from an independent study (*14*). The study consisted of an intersection investigation of 14 problem intersections in which the objectives were to

1. Select appropriate intersections for detailed analysis,
2. Recommend improvements to selected intersections that would alleviate existing problems, and
3. Document the procedures used.

Seven of the 14 intersections investigated were four-way intersections directly comparable with the knowledge already in the expert system. At present, the isolated intersection part of the system contains knowledge relevant only to four-way intersections.

Traffic counts from the study were translated into symbolic form (low, medium, high, and saturated) and fed into the expert system. General information gathered during the intersection study was used as the general (overview) information for the expert system.

In five of the seven intersections, recommendations made by the expert system program closely followed the recommendations made by the consultants. However, recommended actions such as "Check the operation of the vehicle actuation, as it is not working properly at present" were not suggested by the expert system because it did not yet contain any information relevant to this type of problem.

In each case, the TSM expert system provided a broader range of recommended actions than those provided by the consultants. This outcome is to be expected because the human expert is inclined to filter out the less than ideal solutions to a given problem. No such filtering mechanism was built into the expert system. On the other hand, the expert system's approach of giving a full range of solutions could be regarded as an advantage because of its consistency and comprehensiveness. Humans are sometimes inclined to get in a rut and offer only their personal and familiar solutions.

The major difference between the recommendations of the expert system and the consultants' study was that the expert system program's results were qualitative and required further investigation and analysis before a detailed recommendation could be implemented.

CONCLUSION

An expert system was constructed that addresses the generation and analysis of alternative solutions in the TSM process. In particular, the steps taken in structuring and analyzing the TSM transportation knowledge (the knowledge engineering) into a form acceptable to the expert system model were described.

Experience gained in the construction of the system indicates that the key to successful expert system construction is in the knowledge engineering. Without a clear understanding of how an expert goes about solving a problem, an expert system project is unlikely to succeed.

The methodology described has been effective in providing a structure for the difficult task of encapsulating human expertise within a computer program to generate multisolutions to problems within TSM.

Knowledge engineering procedures have had a useful indirect benefit on the TSM process. The formal process of gathering and structuring the TSM knowledge from disparate sources has provided a consistent approach to the classifying of knowledge and information within TSM. In addition, because of the expert system's consistent and rapid response to a given problem situation, it is likely that the completed system will provide excellent training in the TSM methodology.

ACKNOWLEDGMENTS

Some of the concepts discussed in this paper were developed during work meetings and informal discussions with G. V. F. Timmermans, D. G. M. van Riet, H. P. Leifheit, and S. P. Gerber. Their contributions are gratefully acknowledged.

REFERENCES

1. W. Metrey. An Assessment of Tools for Building Large Knowledge Base Systems. *AI Magazine*, Vol. 8, No. 4, Winter 1987.
2. E. C. Chang. Using Expert Systems to Select Traffic Analysis Software. In *Transportation Research Record 1145*, TRB, National Research Council, Washington, D.C., 1987.
3. E. C. Chang. Application of Expert Systems to Left-Turn Signal Treatment. In *Transportation Research Record 1145*, TRB, National Research Council, Washington, D.C., 1987.
4. M. L. Maher, ed. *Expert Systems for Civil Engineers.* ASCE, New York, 1987, pp. 129–143.
5. J. H. Batchelder, M. Golenberg, J. A. Howard, and H. S. Levinson. *NCHRP Report 263: Simplified Procedures for Evaluating Low-Cost TSM Projects—Users Manual.* TRB, National Research Council, Washington, D.C., 1983.
6. A. B. Tomecki. *Proposals of Organizational Procedures for the TSM Process.* Technical report RT/24, CSIR, NITRR, Pretoria, South Africa, 1986.
7. TSM Application in a Local Authority. User Guide 1, Public Communication of Improvement Needs and Potential, Report No. 23/89.
8. KES II, Knowledge Base Authors' Manual. Vol. I, Production Subsystem.
9. KES II, Knowledge Based Authors' Manual. Vol. II, Hypothesis and Test.
10. P. Jackson. *Introduction to Expert Systems.* Addison—Wesley, Reading, Mass., 1986.
11. F. Hayes-Roth, et al. *Building Expert Systems.* Addison—Wesley, Reading, Mass., 1983.
12. S. H. Weiss and C. A. Kulikowski. *A Practical Guide to Designing Expert Systems.* Rowman and Allanheld, Totowa, N.J., 1984.
13. K. Parsaye and M. Chignell. *Expert Systems for Experts.* John Wiley, New York, 1988.
14. Intersection Study Analysis, ORMET-Van Niekerk, Kleyn and Edwards, File No. 8/16, ORMET, Springs Town Council.

Publication of this paper sponsored by Committee on Transportation System Management.