1986

TOWARD INTELLIGENT DECISION SUPPORT SYSTEMS: SURVEY, ASSESSMENT AND DIRECTION

Timothy R. Hill
Indiana University

TOWARD INTELLIGENT DECISION SUPPORT SYSTEMS: SURVEY, ASSESSMENT AND DIRECTION

Timothy R. Hill
Department of Operations and Systems Management
Graduate School of Business
Indiana University

ABSTRACT

A survey of relevant literature serves as the basis for an assessment of research on integration of decision support systems and artificial intelligence. The analysis identifies the need for a unifying framework with which to direct such research. The characteristics required for such a framework are highlighted and shown to be well-suited to the artificial intelligence concept of deep knowledge. A deep knowledge architecture for intelligent decision support systems is presented and proposed as a basis for integration of the two disciplines.

INTRODUCTION

As research into the application of computer technology continues to advance and meet its original challenges, higher expectations arise for the future. After adequately mastering many of the research issues involved in well-structured tasks such as transaction processing, interest has shifted toward more complex, unstructured tasks including the support of decision making.

The concept of a decision support system (DSS) has been developed to describe a computer system which aids the decision maker, ideally as an intelligent consultant, in solving unstructured or semi-structured problems. Research in DSS, however, is still in its infancy, impeded by the lack of any guiding formal theory. Many positional pieces suggest different directions for development and have led to a number of prototypical systems which have achieved only limited success in providing intelligent support. Recently, researchers have looked to the area of artificial intelligence (AI) for techniques which might be used to dramatically enhance the power of DSSs.

The following analysis focuses on the issues involved in incorporating AI concepts into DSS design. The following section, a survey and review of previous assessments of this area, first establishes the current state of development and then identifies the need for a new, unifying approach. In the next section, the recently emerging AI concept of "deep knowledge" is explored and suggested as the key to achieving a truly intelligent DSS. A review of the deep knowledge literature then serves to introduce the concept, preceding the proposal of a deep knowledge structure intelligent decision support systems.

SURVEY AND ASSESSMENT

Previous Assessments

The incorporation of AI concepts into DSSs has been a topic of great interest in the research community. A review and analysis of the current body of research work provides an assessment of the state of development in this area and helps to identify critical issues which have yet to be addressed. Turban and Watkins (1986), present a survey aimed at introducing and discussing the issues involved in integrating DSS and the AI concept of expert systems (ESs). Admittedly, this is a difficult task due to the nature of the research being examined. Indeed, there appears to be little agreement over the appropriate issues or directions of development to be followed by researchers in this area.

In discussing an ES as an example of a DSS, the authors cite opposing viewpoints from the relevant literature. Since the concept of DSS is rather vaguely defined, various authors have argued both that an ES is a DSS and that it is not. ES has been shown by some to fit neatly into DSS categories, while others have suggested that ES fail to fulfill some critical DSS functions.

The authors present a table of differences which distinguish the two types of systems based on a number of criteria. Specifically, they point out that ESs make rather than support decisions, disqualifying them as true DSSs. They further state that ES typically involve a closed-world assumption in which the problem domain is circumscribed, confining the performance of the system. In contrast, the authors suggest that a true DSS must be adaptive and flexible, with the capacity to evolve.

Based on a review of relevant literature, the authors identify and model two general perspectives taken in attempting to integrate ES and DSS. The first model represents the use of multiple ESs, each designed to support a different component of the traditional DSS architecture. The second model views an ES as a new, additional component of the traditional DSS architecture. Several studies based on each perspective are reviewed with the conclusion that the greatest potential is afforded by using different ESs to enhance each component individually.

Two other valuable conclusions arise from the paper. First, it is apparent that the issue of integrating AI research into that of DSS is the basis for legitimate, growing interest in the research community. Several attempts at integration have been and are being pursued, and some limited success seems to have been realized. Secondly, the wide variety of approaches and lack of any apparent accepted direction of development suggests the dire need of some very general unifying concept. This is a prerequisite to the healthy, productive evolution of the discipline, allowing the constructive progression of research. Unfortunately, the model suggested by Turban and Watkins (1986) is directed toward the unstructured, segmented interjection of ES throughout a DSS, and therefore is insufficient for guiding future research.

A major obstacle to integrating these disciplines comes from the difficulty in maintaining familiarity with two widely-varied and quickly-evolving bodies of literature. Almost all of the literature cited by Turban and Watkins (1986) comes from the DSS community, omitting several instances of relevant AI work. Indeed, several of the ES-DSS differences they present have been blurred by recent AI developments (as will be shown later in this paper). For DSS to successfully reap the benefits of AI integration, full attention must be paid to the AI research community.

Hwang (1985) presents a thorough survey and analysis of related research. Although the topic is defined as "automatic model building systems," the author's treatment is broad enough to encompass many DSS design issues and includes a heavy AI influence.

The author begins with a general description of the decision-making process and identifies two issues which must be resolved in order to achieve the ideal case of a DSS, one which acts as an OR/MS consultant:

1. In the absence of any relevant traditional model for analysis, how can the DSS improve the quality of the decision-maker's solution?
2. If such a model does exist, how might the DSS support the decision maker in selecting and constructing the model?

The author claims that traditional OR/MS and DSS approaches are inadequately addressing these issues while the field of AI offers great promise for the solution.

A review of AI concepts concerning knowledge-based systems is also presented, drawing an analogy between the construction of an expert system and the building of an analytical model. In a knowledge-based domain then, model building refers to the knowledge acquisition task. In addition to research on the extraction of knowledge from human experts, research on several alternative knowledge acquisition techniques are discussed, including, for instance, programs which can automatically infer causal relationships from large databases.

On this basis, the author proposes the solution to the previously proposed questions. First, if no analytical model exists, a knowledge-based model may be constructed to capture an expert's knowledge of the problem domain. Secondly, for cases in which an analytical model does exist, the expertise of the OR/MS consultant may be captured in an ES for access by the decision-maker. The author, however, is careful to point out that the extraction of this knowledge on model building *within* a domain may differ considerably from that for knowledge *about* a domain. Apparently, this task has proven to be very difficult in previous research which has met with only limited success. Such knowledge is said to be much more difficult to acquire.

Following the survey, the author proposes the concept of an intelligent decision support system (IDSS) as a goal for DSS research. Such a system would support the decision maker throughout the entire decision-making process, acting somewhat like a human decision analyst. An IDSS, therefore, is characterized by the author as being able to:

1. analyze the problem and associate it with a solution approach,
2. construct or search for appropriate models based on the solution approach,
3. execute the model to obtain solutions, and
4. interpret the solution and document the lessons learned.

The author suggests that DSS research has not addressed the first two of these issues, but rather has concentrated on the third. The fourth is said to be a research topic of current interest. Deficiencies are pointed out in some of the surveyed research which attempts to support the entire decision-making process and the ability to construct new models is identified as the key. The author echoes the argument raised by Turban and Watkins (1986), saying that most existing intelligent systems perform well only in a very limited domain, insufficient for supporting the wide variety of decision activities addressed in an IDSS.

As does the Turban and Watkins article (1986), this paper suffers from the lack of any useful unifying concept. Though the categorization of the research is helpful, the definition of the categories is still somewhat weak. Of course, this problem is inherent in the widely varied nature of the research being studied.

In order to answer the need for some unifying concept to direct the research toward IDSS, the characteristics required of such a framework must first be identified. A brief review and analysis of a sample of three relevant research projects will serve to clarify the issues and point to a logical solution. Specifically,

1. support will be derived for the importance of the IDSS, as suggested by Hwang (1985), and for the feasibility of IDSS support of the entire decision process,
2. evidence will be compiled to show that the differences between ES and DSS, which were proposed by Turban and Watkins (1986), are being obscured by recent AI developments, legitimizing the use of ES as a DSS, and
3. the need for a unifying framework will be demonstrated by showing the difficulty in integrating the diverse developments and the failure to build on related research.

Analysis of Examples

The first two examples of research are attempts to explicitly link AI and DSS. While the references cited by Turban and Watkins (1986) were primarily DSS-based, most of the references by Hwang (1985) are to instances of AI research, linked to DSS only by their impact as suggested by the author. In contrast, the following two examples represent concrete research effort aimed at the integration of disciplines.

In Duda and Reboh (1983), the authors draw upon their experience in developing an extensive expert system to gain insight into decision making. It is the proposition of the authors that their knowledge base approach to a mineral exploration advisor is an advance in modeling human decision making. Clearly, they are taking the view described in Turban and Watkins (1986) that this ES is supporting the decision-making process and is thus a DSS.

The emphasis of this research is on the development of the knowledge base. The main advantage is the allowance for uncertainty in the knowledge base. Beliefs may be described and incorporated. And in the inference process, evidence is gathered in support of alternative solutions.

This paper has major implications for DSS-AI integration research. First, it supports Hwang's (1985) concept of IDSS by advancing the state of model building capabilities. It is a step toward better models of domain knowledge, and toward achieving the ability to model the expertise of the OR/MS consultant.

Secondly, it refutes much of the distinction proposed by Turban and Watkins (1986) between ES and DSS. For instance, the objective of this system is clearly to assist humans (like DSS) rather than to replicate and replace a human as suggested by Turban and Watkins (1986) for typical ES. In addition, the system seems to fit into the DSS side of their table, in the category of "major orientation," with a decision making slant as opposed to "transfer of expertise." Furthermore, this system merely gathers and reports the evidence supporting the various alternatives under consideration. It is the human who makes the final decision. This point would also land the system in the DSS categorization of Turban and Watkins (1986). These facts suggest that knowledge base systems might indeed develop all of the capabilities required to fully support decisions. Certainly, they should not be ignored in the pursuit of the goal of IDSS.

Nakamura, *et al.* (1982) provide a second example of research which explicitly links AI and decision support. In this case, however, there is an entirely different emphasis. The reported system draws on a knowledge base which is extracted from textual documents within a particular domain through the use of cognitive maps.

The system aids the user in dealing with the complexity of recent societal problems for which a large number of causal relationships exist, often spanning multiple disciplines. The described DSS includes three retrieval modes which allow the decision maker to extract different forms of knowledge easily, supporting the manipulation and analysis of cognitive maps.

Viewing this paper with respect to the DSS-AI integration issues previously raised, the characteristics of a prescriptive strategy may be more clearly articulated. Unlike the previously discussed paper which dealt with a very narrowly defined knowledge domain, this one represents an attempt to traverse multiple domains and address more general problems. Such work brings a new dimension to the area of DSS-AI integration and is especially important in light of the argument made by Turban and Watkins (1986). They highlighted the closed-world assumption as a fundamentally constraining characteristic of ES which limited its ability to fulfill a DSS role. This concern was echoed by Hwang (1985) in citing the typically narrow domain of ES. Nakamura, *et al.* (1982), however, take a step toward answering this issue and advancing the evolution of IDSS.

The intuitive appeal of this work, however, emphasizes the need for a general framework for development. Like the paper by Duda and Reboh (1983), it seems certain that this research has valuable implications for AI-DSS research. Yet it is very difficult to place either one into any meaningful overall scheme. It is thus nearly impossible to combine the benefits of the two pieces or even build upon the research of either one. The need for a framework is further demonstrated by a review of the following paper.

Unlike the previous two studies, work by Michalski (1980) makes no reference to decision support. Yet, intuitively it appears to provide another development of importance for DSS-AI integration. This work addresses the problem of grouping items by some set of criteria into categories called clusters. The author develops and refines the supporting theory and presents an algorithm for automatically "learning" concepts based on the clustering of data. Clearly, this work is related to the category described by Hwang (1985) which included methods of inferring knowledge from the computational analysis of large amounts of data.

Unfortunately, this work underscores the need for a unifying framework for DSS-AI integration. It is written strictly within the realm of AI/computer science and is particularly difficult to relate to DSS. The theoretical development is quite rigorous and the examples are very brief and concise. The paper is clear enough, however, to assure its value to IDSS. A useful general framework for integrating DSS and AI must be broad enough so that even basic research such as this, coming from the far end of the spectrum, may be meaningfully incorporated.

In the final example of related research, Swartout (1983) presents yet another development of interest to DSS. Like the previous paper, however, there is no explicit reference to decision making issues.

The author describes a shell system for constructing ESs. The central theme is the intelligent explanation of the reasoning process by the generated ES. The trust and confidence of the user are assumed to be enhanced by the ability to examine the ES's logic. The shell system interacts with an expert, first extracting most general fundamentals of the domain. The interaction is then controlled by the shell system in order to fully develop the implications of those principles. When the resultant ES queries the user for information, the user may respond with a question, "WHY?" The ES then examines the refinement structure stored in the knowledge base and reports the reasoning used.

The examples indicate impressive results in which the explanations appear to be intelligently formed. Indeed, they appear to be highly flexible and constructed differently, depending on the particular prior circumstances which have occurred to that point in the consultation session. The system thus achieves a degree of two-way interaction between the system and the user, representing more of a partnership in the reasoning process.

The review of relevant literature provides strong implications for DSS-AI integration research. Their summary and analysis will serve as a basis for the suggestion of a unifying framework for future development. First, the evidence indicates that the concept of IDSS, as proposed by Hwang (1985), has practical feasibility and is thus worthy of research effort. Indeed, several features of such a system have shown much promise already, although most have been developed in isolation from the others.

Furthermore, in accordance with Hwang's (1985) description of an IDSS, many of the phases of decision making are being supported with some degree of intelligence. For example, the automatic development of models of knowledge is advanced by the work of Michalski (1980) in developing concepts through the clustering of data. Interpretation of the results of the inference process is enhanced by the ability to automatically generate explanations of reasoning, as in Swartout (1983). The knowledge base of cognitive maps presented in Nakamura, *et al.* (1982) aids the association of the problem to a solution approach through clarification of causal links. Such research represents the first primitive steps toward IDSS and testifies to its feasibility.

The research further clarifies the issue concerning the appropriateness of a knowledge base system for IDSS. Many of the ES-DSS differences presented by Turban and Watkins (1986) are being resolved by recent AI developments. For example, Swartout (1983) reported a technique to support two-way interaction in an ES, approaching the "consultant" concept so important in DSS. Duda and Reboh (1983) described a knowledge based system in which the objective was to assist humans in making decisions. This practical approach does not conform to the goal of replicating or replacing humans as suggested by Turban and Watkins (1986). Furthermore, the ES described by Duda and Reboh (1983) handled complex numerical manipulation in addition to that of symbolic data. This feature spans both the DSS and ES descriptions of manipulation as depicted in Turban and Watkins (1986).

Their delineation by "problem area," however, remains an issue of concern. Here, they differentiate ES as being confined to a narrow domain, an inappropriate restriction for DSS. Their opinion that DSS must maintain flexibility and be unbounded by a "closed-world assumption" is echoed by Lee (1983) and Hwang (1985). The research reported by Nakamura (1982) is an early attempt at answering this concern by bridging multiple domains. However, this seems to be the fundamental problem remaining to be answered in a knowledge base IDSS.

Finally, the review underscores the need of a unifying framework for DSS-AI research development. Although many of the studies indicate great potential, they are generally performed in isolation from each other. The wide diversity of approaches makes it very difficult to integrate the small steps into any major advancement. The work of Nakamura, *et al.*

(1982) and Duda and Reboh (1983) seem to overlap each other in producing two nearly complete systems, while the conceptual clustering work of Michalski (1980) is so basic and theoretical that, on the surface, it seems nearly unrelated to the others. In order to maintain a logical and efficient progression, future research requires some overall structure, general enough to integrate the diversity of advancements.

In summary, a unifying framework is essential for the future development of AI-DSS research. The framework must provide the generality to incorporate a wide variety of relevant advancements. Furthermore, it should be directed toward the evolution of the IDSS as suggested by Hwang (1985). Knowledge based systems possess many of the characteristics required for decision support. A knowledge based system structure might thus serve as the basis for IDSS development, and therefore as a framework for DSS-AI research. However, such a structure must provide flexibility and not be constrained to a narrow domain. In the next section, the recently developed AI concept of deep knowledge is introduced and suggested as an approach to resolving these difficult issues.

DIRECTION

Deep Knowledge

The recently developed AI concept of deep knowledge (DK) appears to hold great potential for the advancement of research toward IDSS. Indeed, the following analysis will propose a key role for DK in a unifying structure for AI-DSS integration research. First, a review of two seminal DK papers will serve to introduce the subject. The utility of the approach will then be demonstrated by examining two projects which have realized successful implementation of DK systems.

In an AI direction piece by Hart (1982), the concept of deep knowledge was first articulated. The author focused on conceptual complexity as one of the most important questions facing AI. He began by drawing a distinction between surface and deep systems. Surface systems are described as "those having no underlying representation of such fundamental concepts as causality, intent, or basic physical principles." In contrast, these features are said to be present in deep systems. Several brief examples of actual systems are cited to clarify the delineation.

According to the author, deep systems involve greater conceptual complexity and, therefore, greater startup costs. However, they should be capable of solving tasks of greater difficulty than a surface system in the same domain. This relationship is clarified in a depiction borrowed by the author from a more limited discussion by Gary Hendrix (see Figure 1).

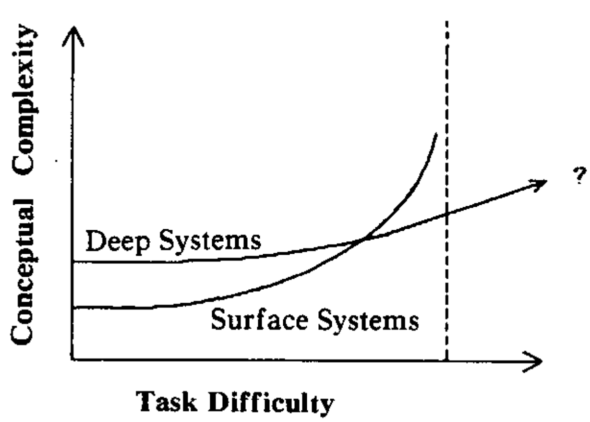

Figure 1: Deep vs. Surface Systems.

Thus it is suggested that either type of system might be able to solve problems below some threshold level of difficulty. In fact, surface systems would perform adequately over much of this range with a lower level of conceptual complexity and, therefore, with lower startup expense. However, as the threshold is approached, the complexity of a surface system grows rapidly, surpassing that of a deep system. This is a reflection of the excessive patching and rigging required to solve more difficult problems without knowledge of underlying principles. At some point, the task difficulty is simply beyond the ability of the surface system.

Deep systems, however, have the potential to exceed the threshold and deal with possibly unlimited task difficulty. Thus, a system which contains rules about tests, symptoms, and treatments of lung disease might be adequate for many cases. However, one which is built with knowledge of the principles of the human respiratory system would be able to deal with the most difficult cases, perhaps even those which were unforeseen during the design phase of the system.

The author goes on to extend the discussion to multi-level systems. Since different levels of knowledge may be suited to different tasks, it is suggested that a system might include more than one level of knowledge. A hypothetical system for petroleum reservoir engineering is examined in which there are three levels of knowledge representations. A surface level model contains simple if-then rules which provide useful information but do not allow the derivation of any causal knowledge. A second level, intermediate model captures some gross notions of causality in the form of volume balance equations. Finally, the most detailed representation of the knowledge involves a set of partial differential equations which are simulated computationally as difference equations. This deep level knowledge embodies the most basic physical laws which govern the problem domain.

The discussion is followed by the articulation of issues in the design of multi-level systems. It is suggested that this concept holds potential for purposes of validation, explanation, and control. The possibilities of non-hierarchical arrangements of levels are also mentioned. Further noted are problems in the representation formalisms for different levels, including the need for inter-level communication. Finally, the author raises the question of how to exploit multi-level designs to realize practical benefits.

The appropriate use of multi-level design is addressed briefly by Michie (1982) in another early paper on DK. Referring to deep systems as "high-road programs" and surface systems as "low-road programs," the author suggests that the DK approach is not always appropriate. Defending the utility of surface systems, he states that most deep systems have produced solutions which are "opaque to the user and unbelievably costly at run time." Reference to research on a multi-level system is then made, highlighting the sensibility of trying first to solve the problem with surface knowledge and resorting to deep knowledge only when necessary.

Davis (1983) reports success with a DK system for electronic troubleshooting. This is an excellent paper which provides great detail and makes several valuable points. The author explains the value of using first principles (DK) in solving problems which are beyond the limits of traditional techniques. Specifically, the objective was to build a system which would capture the skill of an engineer in diagnosing problems in an electronic hardware device which he has never seen. In such cases, the engineer must analyze the problem, based only on his general knowledge and the device schematics. The author argues for the use of primal models of causal interaction to achieve this goal. The system demonstrates the successful implementation of a multi-level approach, allowing simplifying assumptions for a problem to be made initially, and surrendering or retracting them when more complex hypotheses are required. An example is presented to demonstrate this capability before drawing conclusions and implications from the experience. This work clearly supports the use of multi-level DK, both in terms of power and practicality.

A further development in this area is reported by Chandrasekaran and Mittal (1983). Here the authors discuss the use of "compiled knowledge" (CK) as an alternative to DK in diagnostic knowledge based systems. CK is described as being derived from the deeper knowledge in a domain, with the objective of solving a specific set of problems. Lying thus somewhere between deep and surface knowledge, CK is shown to be sufficient for solving many problems without imposing the excessive computational costs associated with DK. The authors note, however, that the CK for a given set of problems may not be sufficient for solving other problems which could theoretically be solved with its underlying DK.

These propositions are consistent with the previously described work of both Hart (1982) and Michie (1982). From the depiction of conceptual complexity versus task difficulty which appears in Hart (1982), it is implied that higher level knowledge is beneficial up to a certain point of difficulty. Beyond that point, however, it is limited, unlike the potential of DK. Michie (1982) suggested that DK alone had limited practical applicability and described the use of multiple levels. This is another issue addressed by Chandrasekaran and Mittal (1983). Clearly, the concept of DK and multi-level knowledge based systems is developing a sound base of research and practical use.

A Unifying Structure for IDSS

Based on the preceding review and analysis of both the DK-related literature and that related to DSS-AI integration, the DK concept may be shown to play a key role in advancement toward IDSS. Certainly, DK represents another AI development which should be incorporated into DSS-AI systems to enhance the power of knowledge bases. But in a more general sense, a DK approach to an IDSS structure may provide a basis for a unifying framework to direct DSS-AI integration research. It is therefore proposed that a DK knowledge based system architecture, directed at IDSS development, serve as the instrument of direction needed to guide research aimed at AI-DSS integration.

From the review of AI-DSS integration research, it was concluded that a knowledge based system might indeed be capable of serving as a basis for decision support if the problem of the narrow domain or closed-world assumption for such a system could be overcome. Both Turban and Watkins (1986) and Hwang (1985) highlighted this as a prominent obstacle to intelligent decision support. Extension of the DK approach, however, suggests a solution to the closed-world problem.

Theoretically, this issue may be resolved by taking the concept of DK to a greater level of abstraction, applying the DK principle to the general area decision making rather than to a specific problem domain. As Figure 1 suggests, the resulting system would not be limited by the closed world assumption, but would be capable of dealing with tasks which were beyond the scope of the original design. By reasoning from the basic fundamentals of decision making, such a system could intelligently formulate "original" solution strategies, extending the boundaries of decision support beyond previous limits of complexity to problems of even less structure.

On this basis a knowledge based structure for IDSS may be proposed, applying the DK concept comprehensively throughout an hierarchy of supporting components. Such a structure is depicted in Figure 2. The multi-level design allows the benefits of DK to be realized throughout the decision making process. For example, DK might also be applied at an abstract level to capture the expertise of the OR/MS consultant, proposed by Hwang (1985) as a critically important but very difficult step toward IDSS. Most importantly, the structure can serve as a vehicle by which IDSS may be allowed to evolve incrementally by incorporating related research into a unifying framework.

The key to such a structure is the application of DK within a multi-level system built on the deepest possible foundation. Thus, the kernel of the system would include the DK, or first principles, of decision making. This knowledge would describe the most fundamental issues involved in formulation and evaluation of alternatives and information theory. This kernel would then drive the system, communicating with the other, higher levels of knowledge in order to achieve its goals.

At the next higher level, the DK of learning and knowledge acquisition techniques would be made available to the kernel. Thus, the best approach to gathering the necessary information would be used to access still higher levels of knowledge. In some cases, the knowledge could be acquired directly from the relevant domain-specific knowledge base. Alternatively, it may be necessary to call upon the OR/MS knowledge to direct an analysis in order to synthesize the desired information.

The knowledge of OR/MS techniques would capture the skill of the OR/MS consultant using the approach presented by Davis (1983). It would have access to the model base and to the domain-specific knowledge bases, which, in turn, would have access to the database. Based on the fundamental principles of OR/MS, this component would enable the system to intelligently select from available models, or to devise unstructured models of knowledge for more complex problems as suggested by Hwang (1985). Furthermore, this knowledge would be used to provide a liaison between the decision maker and the model base, providing supportive insight and expertise.

Similarly, multiple domain-specific DK bases would insulate the user from the raw data of the database. Instead, the knowledge reflected in the data would be available, due to the DK approach, and would be extensible based on basic principles derived. The accumulation of this knowledge might be performed automatically based on the DK base of learning and knowledge acquisition.

In this way, the traditional components of DSS are supported by an underlying intelligence which has the flexibility to learn about and deal

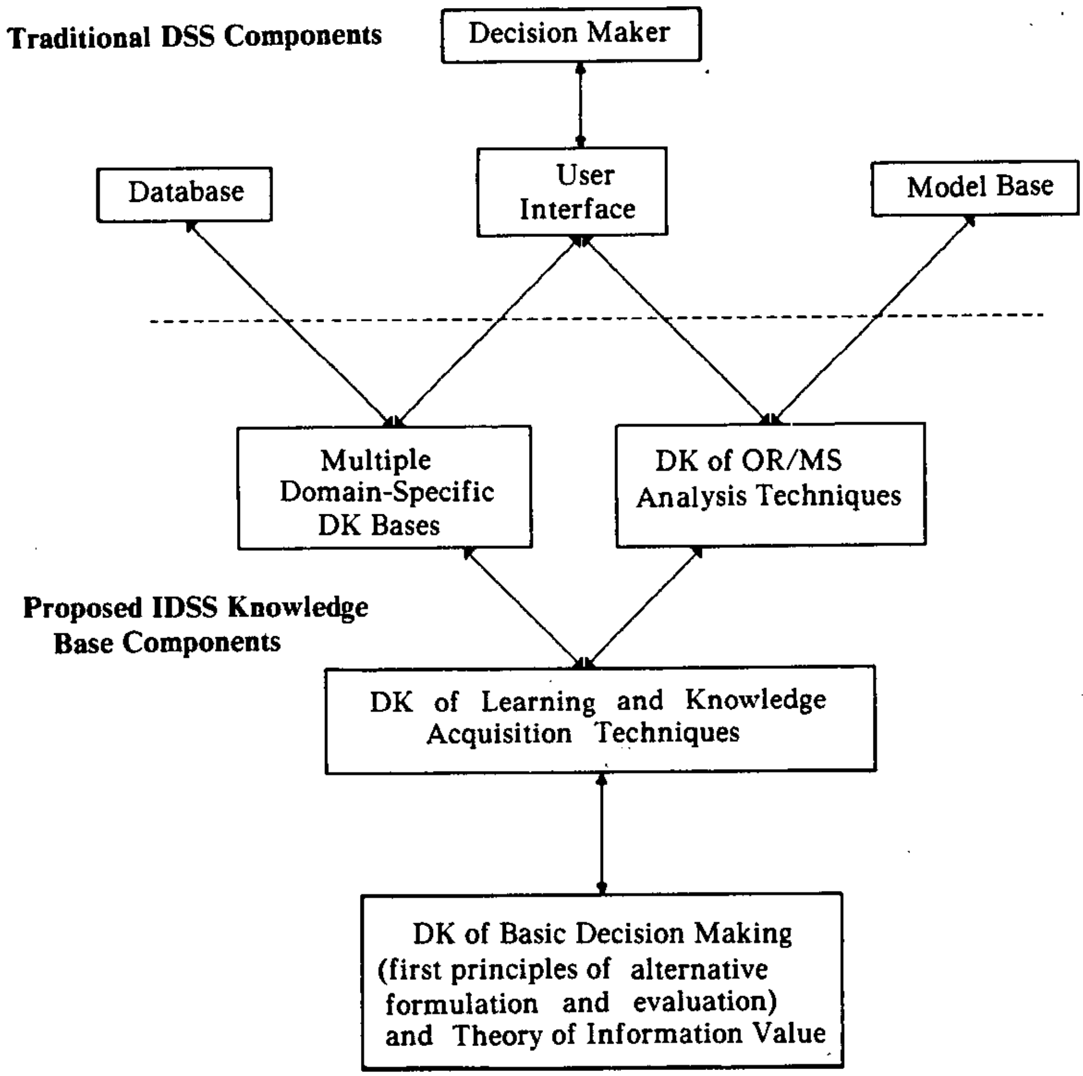

Figure 2: Proposed IDSS Structure.

with problems which are unforeseen in the original design process. Of course, the enormous scope of such a system casts doubt as to its practical feasibility. But for sets of problems which are limited, the compiled knowledge approach outlined by Chandrasekaran and Mittal (1983), might be employed to increase operating efficiency. Furthermore, the modular, hierarchical structure delimits sub-topical areas, constraining the scope of related research projects and providing a framework by which they may be constructively directed.

The sample of relevant research helps demonstrate the use of the IDSS structure to integrate current developments and direct future ones toward a meaningful goal. Michalski's (1980) technique for automatic learning through conceptual clustering would clearly be incorporated into the DK base of knowledge acquisition. It could be called upon during a consultation session to develop new knowledge to support a decision. Or it might run in the background, constantly analyzing the database to build the DK bases of domain-specific knowledge.

Similarly, the domain-bridging system of Nakamura, *et al.*, (1982) appears to have potential impact on the DK base of learning and knowledge acquisition. It might be used as a technique to link the multiple domain-specific DK bases, drawing conclusions by integrating fundamental knowledge from different areas. Of course, the use of these techniques and others would be governed by the deeper, fundamental knowledge of learning and knowledge acquisition.

The work of Duda and Reboh (1983) suggests the incorporation of uncertainty throughout the entire knowledge base, supported by the DK base on decision making and information theory. Finally, the multi-level use of intelligent explanation, as described by Swartout (1983), could be built into the hierarchical structure. The resulting system might meet the call of Turban and Watkins (1986) for a move from "what if" to "why" capabilities, and a transformation from a passive to an active role in the decision-making process.

It is hoped that the IDSS structure may serve not only as a framework by which to organize past developments, but also as a means of constructively directing future research. The global structure of the IDSS may help provide a meaningful goal toward which apparently divergent lines of research may be focused.

SUMMARY AND CONCLUSIONS

Through the review and analysis of literature related to the integration of DSS-AI literature, several conclusions were apparent. First, the relevance of the discipline was established. Several AI developments were shown to promise great beneficial impact for DSS. Secondly the analysis identified the need for a unifying framework to guide such research. The characteristics required of such a framework were also developed. Thirdly, support was shown for the appropriateness of a knowledge based system structure as the guiding framework. The concept of IDSS as proposed by Hwang (1985) was suggested as the basis for such a structure.

The recently developed AI concept of deep knowledge was then examined. A review of four related papers introduced the concept and demonstrated its successful application. Deep knowledge refers to the use of fundamental principles to solve tasks of greater difficulty than previously possible. This concept was suggested as the key to a general IDSS structure. A general structure for IDSS was then proposed as the unifying framework necessary to provide direction for research aimed at DSS-AI integration. The kernel of the system drives the process using the basic principles of decision making. Information needs are communicated from the kernel to the knowledge base containing the fundamental concepts of learning and knowledge acquisition. This component has access to the domain-specific deep knowledge, and to such knowledge of OR/MS analysis. These components are then linked to the traditional components of DSS.

The proposed structure for a knowledge base IDSS provides a scheme into which related research may be placed. In this way, AI developments may be incorporated into the effort to achieve the goal of IDSS. Future research should be directed toward developing the components of this structure and facilitating their interaction. In order to insure the utility of the approach, the kernel of the system should receive first priority. This difficult task will require the identification of the basic principles of decision making and the incorporation of information theory. Through this approach, the evolution of IDSS may be more efficiently achieved.

REFERENCES

Chandrasekaran, B. and Mittal, S. "Deep Versus Compiled Knowledge Approaches to Diagnostic Problem Solving," *International Journal of Man-Machine Studies,* Volume 19, 1983, pp. 425-436.

Davis, R. "Reasoning From First Principles in Electronic Troubleshooting," *International Journal of Man-Machine Studies,* Volume 19, 1983, pp. 403-423.

Duda, R. O. and Reboh, R. "AI and Decision Making: The PROSPECTOR Experience," in *Artificial Intelligence Applications for Business,* Walter Reitman (ed.), Ablex Publishing Corp., Norwood, New Jersey, 1983.

Hart, P. E. "Directions for AI in the Eighties," *SIGART Newsletter,* Volume 79, January 1982.

Hwang, S. "Automatic Model Building Systems: A Survey," *Proceedings of the 1985 DSS Conference,* San Francisco, California, April 1985, pp. 22-32.

Lee, R. M. "Epistemological Aspects of Knowledge Base Decision Support Systems," in *Processes and Tools for Decision Support,* by H. G. Sol (ed.), North-Holland Publishing Co., Amsterdam, Holland, 1983.

Michalski, R. *Knowledge Acquisition Through Conceptual Clustering: A Theoretical Framework and Algorithm for Partitioning Data into Conjunctive Concepts,* University of Illinois Press, Urbana, Illinois, 1980.

Michie, D. "High-road and Low-road Programs," *AI Magazine,* Volume 3, Number 1, 1982.

Nakamura, K., Sosuke, I. and Sawargi, T. "Decision Support Using Causation Knowledge Base," *IEEE Transactions on Systems, Man, and Cybernetics,* Volume SMC-12, Number 6, November/December 1982.

Swartout, W. R. "XPLAIN: A System for Creating and Explaining Expert Consulting Programs," *Artificial Intelligence,* Volume 21, 1983.

Turban, E. and Watkins, P. R. "Integrating Expert Systems and Decision Support Systems," *MIS Quarterly,* Volume 10, Number 2, June 1986, pp. 121-136.