

DTIC FILE COPY

1

Research Note 88-41

AD-A197 386

COMPUTER BASED INTELLIGENCE SUPPORT:
AN INTEGRATED EXPERT SYSTEM AND
DECISION SUPPORT SYSTEMS APPROACH

Andrew P. Sage
School of Information Technology and Engineering
George Mason University

Adolfo Lagomasino
AT&T Bell Laboratories

for

Contracting Officer's Representative
Michael Drillings

ARI Scientific Coordination Office. London
Milton D. Katz, Chief

BASIC RESEARCH OFFICE
Michael Kaplan, Director



U. S. Army

Research Institute for the Behavioral and Social Sciences

June 1988

Approved for public release; distribution unlimited

DTIC
ELECTE
JUL 08 1988
S H D

88 7 06 078

U. S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES

A Field Operating Agency under the Jurisdiction of the
Deputy Chief of Staff for Personnel

EDGAR M. JOHNSON
Technical Director

WM. DARRYL HENDERSON
COL, IN
Commanding

Research accomplished under contract
for the Department of the Army

George Mason University

Technical review by

Dan Ragland

This report has been cleared for release to the Defense Technical Information Center (DTIC). It has been given no other primary distribution and will be available to requestors only through DTIC or other reference services such as the National Technical Information Service (NTIS). The views, opinions, and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other official documentation.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ARI Research Note 88-41	2. GOVT ACCESSION NO. --	3. RECIPIENT'S CATALOG NUMBER --
4. TITLE (and Subtitle) Computer Based Intelligence Support: An Integrated Expert Systems and Decision Support Systems Approach		5. TYPE OF REPORT & PERIOD COVERED Interim Report October 84 - October 85
		6. PERFORMING ORG. REPORT NUMBER --
7. AUTHOR(s) Andrew P. Sage and Adolfo Lagomasino		8. CONTRACT OR GRANT NUMBER(s) MDA903-82-C-0124
9. PERFORMING ORGANIZATION NAME AND ADDRESS School of Information Technology and Engineering, George Mason University Fairfax, VA 22030		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 2Q161102B74F
11. CONTROLLING OFFICE NAME AND ADDRESS U.S. Army Research Institute for the Behavioral and Social Sciences, 5001 Eisenhower Avenue, Alexandria, VA 22333-5600		12. REPORT DATE June 1988
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) --		13. NUMBER OF PAGES 29
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE --
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) --		
18. SUPPLEMENTARY NOTES Michael Drillings, contracting officer's representative		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Knowledge Representation, Expert Systems, Artificial Intelligence, Intelligence. (ST...) Cognitive Science,		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This research note summarizes research and design needs in the area of computer based intelligence support, particularly the command and control area. It is argued that an integrating approach is needed. A discussion of knowledge representation emphasizes this, and an approach that accomodates both probabilistic and logical support, and which is able to cope with several types of imperfect information is described. 100, 122		

UNCLASSIFIED

**Computer Based Intelligence Support:
An Integrated Expert Systems and Decision Support Systems Approach**

by

ANDREW P. SAGE

**School of Information Technology and Engineering
George Mason University
Fairfax VA 22030**

ADOLFO LAGOMASINO

**AT&T Bell Laboratories
Holmdel, NJ**

This paper summarizes research and design needs for intelligence based support, especially in areas of command and control. It argues strongly for an approach that integrates, especially in the areas of inference analysis, contemporary approaches in artificial intelligence for expert system construction and management science or systems engineering approaches for the design of decision support systems. A discussion of knowledge representation and information processing highlights needs for this. An approach that accommodates both probabilistic and logical support, and which is able to cope with several types of imperfect information is described.

first draft: November 11, 1984

revised: October 20, 1985

in Expert Systems for Managers, B. G. Silverman (ED), TIMS series in Artificial Intelligence and Management Science, Addison Wesley.

The research described here was supported by the U. S. Army Research Institute under Contract MDA 903-82-C-0124.

1. Introduction

Intelligent systems for decision support are computer implemented procedures that seek to combine expert knowledge about a domain with expert methods of conceptualizing and reasoning about that domain. They integrate this with "formal" methods of reasoning about the domain. The inferential power of such computer programs rests upon a knowledge base that puts together factual information about the domain with the heuristics or informal "rules of thumb" experts use to rapidly find solutions to problems and with the formal reasoning methods that are needed when approaching an unstructured problem about which experiential familiarity is slight.

The goal of an intelligent system for decision support is to encode in a computer program the facts an expert has and the methods of reasoning about them, together with formal methods of reasoning about unstructured situations. In that sense, an intelligent system may be viewed as a descriptive model of an expert reasoning process about a problem domain and a normative system to aid in formal reasoning when this expertise does not exist. Although intelligence support systems have the potential of encoding cognitive biases or prejudice, they may be very attractive to planners and decision makers because they could be set up as on-line decision support systems in situations where time limitations are an important factor. There are methods, currently under investigation, that enable detection of inconsistent reasoning and it is hoped that such efforts may be extended to enable detection of, at least some, of the identified cognitive biases associated with human information processing.

2. Knowledge Representation and Processing in Decision Support Systems Design

Approaches that will enable effective knowledge representation, and associated inference activities, in large knowledge bases have been the subject of investigation for many researchers. An appropriate representation can be used to describe the four different types of factual knowledge elements that may be captured in a knowledge base: objects, events, performance, and metaknowledge. It will assist in identification

Codes

and/or
Special

List



A-1

of the values that need to be associated with facts in order to enable judgment and choice. The purpose of a particular knowledge representation is to enable the use of knowledge for: retrieving factual information from the knowledge base that is judged relevant to the task at hand, reasoning about these facts in the search for a resolution of the task requirements, and acquiring more knowledge. Several approaches for the representation of knowledge as they apply to the control and generation of dialogue for human-system interaction are discussed in Sage and Lagomasino [1984].

2.1 Knowledge Base Management System [KBMS]

A knowledge base management system is one of the three fundamental components of an intelligent system for decision support. In almost every instance in which there will exist multiple decisionmakers, there will exist the need for individual knowledge bases and organizational knowledge bases. Some of the desirable characteristics of a DBMS include the ability to cope with a variety of data structures that allow for probabilistic, incomplete, imprecise, and other forms of imperfect data; and data that is unofficial, and personal, as contrasted with official and organizational.

In order to construct a knowledge base, it is first necessary to identify a knowledge representation scheme. A representation scheme is a collection of data structures, operations that may be applied on the data structures, and integrity rules which are used to constrain or otherwise define permissible values of the data. There are at least five models that may be used to represent data. The most elementary of these is the individual record model. The relational model is a powerful generalization of the record model. A relation is the fundamental data structure in the relational model, and there may be a number of fields in any given relation. The relational model enables mathematical set operations in terms of addition of new records, updating fields within existing records, creating relations that may be contained in records, deleting relations that may be contained in records, joining or combining two or more relations based on their containing common fields, selecting

records by virtue of their containing certain specified relations, and projection such as to enable selection of a subset of the fields that exist in a relation.

The hierarchical or tree data model is a relatively efficient representation of data. In a hierarchical model, the structure represents the information that is contained in the fields of a relational model. In a hierarchical model, there will be certain records which must exist before other records can exist since every data structure must have a root record. Because of this structured aspect of the model, it will be necessary to repeat some of the data that need be stored only once for a relational model. The network model is a generalization of the hierarchical model in that there are links between records which enables a given record to participate in several relationships. There are often major problems associated with insertions, deletions, and updating in both the hierarchical and network data models due to the need to maintain a consistent data base. These do not exist in the relational model since the same data is never entered more than once. Also there is additional search complexity since a search can start anywhere in the network structure. Searches are, however, generally more efficient than they are in the relational model.

Due to the need to accommodate intelligent capabilities in a decision support system, it is desirable to consider a production rule model as a fifth data model. This will enable inferences to be made. Thus this is a particularly desirable form of data model when we desire to use many predictive management information system capabilities. The "if then" type response to "what if" queries is especially natural in this representation. As will be shown, this representation allows for efficient meta-level control, through decision support type approaches, of the production rules in an expert system knowledge base. Knowledge representation formats also include frame, script, and schema representations. These representations are much like the cognitive maps that humans construct of the world around them. At this point, there does not seem to exist operational data models based on these representations, and surely this is an area of contemporary research need. Additional discussion of data

base management management system [DBMS] design approaches can be found in Sprague and Carlson [1982], and Date [1977, 1983]. The three volume Handbook of Artificial Intelligence edited by Barr, Cohen and Feigenbaum [1981,1982] and the text by Sowa [1984] discuss knowledge representation from an AI perspective.

2.2 Model Base Management System [MBMS]

It is necessary to provide inference capability in an intelligent system. This requires some sort of model base management system [MBMS]. It is through the use of model base management systems that we are able to provide for sophisticated analysis and interpretation capability in a decision support system. The single most important characteristic of a model base management system is that it should enable the decisionmaker to explore the decision situation through use of the knowledge base by a model base or algorithmic procedures and management protocols. This can occur through the use of modeling statements, in some procedural or nonprocedural language, through the use of the model subroutines, such as mathematical programming packages, that are called by a management function, and through the use of data abstraction models. This approach facilitates updating and use of the model for explanatory and explication purposes.

The use of multiple models can potentially accommodate the desire of the typical decisionmaker for flexibility. Thus a mixed scanning approach might be incorporated in which a conjunctive or disjunctive scanning mechanism is used to allow for an individual scan to eliminate grossly unacceptable alternatives. After this is accomplished, further evaluation of alternatives might be accomplished by a compensatory tradeoff evaluation, or one based on a dominance search procedure [Sage and White, 1984].

To provide flexibility, the MBMS should provide, upon system user request, a variety of prestructured models that have been found useful in the past, such as linear programming and multiattribute decision analysis model, and procedures to use these models. It should also allow for the development of user built models and heuristics that are developed from established

models that will either become permanent parts of the MBMS or which will be considered as ad hoc models. It should also be possible to perform sensitivity tests of model outputs, and to run models with a range of data to obtain the response to a variety of what if type questions.

2.3 Dialogue Generation and Management System [DGMS]

The dialogue generation and management system portion of a decision support system is designed to satisfy knowledge representation, and control and interface requirements of the intelligent system for decision support. It is the DGMS that is responsible for presentation of the outputs of the system to the decisionmaker and for determining, acquiring, and transmitting their inputs to the KBMS and the MBMS based on its knowledge about the decisionmaker's goals about the specific decision situation. Thus the DGMS is responsible for producing output representations, for obtaining the decisionmaker inputs that result in the operations on the representations, for interfacing to the memory aids, and for explicit provision of the control mechanism that enable effective dialogue between the user and the KBMS and MBMS.

There are a number of possible dialogues. These are inherently linked to the representational forms that are used for the DBMS and MBMS. Menus, spreadsheets, tradeoff graphs, and production rules are some of the formats that may be used as a basis for dialogue system design. Generally, several of these should be used as the support system user may wish to shift among these formats as the nature of issues and experiential familiarity with the issues changes. The DGMS should be sufficiently flexible such as to allow review and sensitivity analysis of past judgments, and to provide partial judgments based upon incomplete information. Of course, the DGMS should be "user friendly" through provision of various HELP facilities, prompting the decisionmaker, and other abilities that supports the knowledge of the support system user.

3. Management of Intelligent Systems

Limits associated with the cognitive capacity of the human mind, time limitations, and many other competing concerns of the decisionmaker are constraints that affect adequate formulation, analysis, and interpretation of complex large scale issues. A design goal for intelligent systems and processes that will assist in various problem solving tasks is to reduce, to the extent possible, the effects of the aforementioned constraints so as to enable the efficient and effective use of information that will lead to quality judgments in routine and familiar task environments and in unfamiliar task conditions.

An appropriate framework in which knowledge could be organized and utilized efficiently and effectively is desired. This is especially needed as studies have shown that the way in which a task is framed exerts a very strong influence upon the way in which task requirements and task resolution efforts are determined [Kahneman, Slovic and Tversky, 1982; Sage, 1981]. This requires that we be able to address the modeling of intelligent systems for decision support from several perspectives. Our interest here is to describe implications that arise in the design of intelligent systems for decision support that incorporates systems management activities. Of particular interest will be those components at the interface between the cognitive process level of systems management and the problem level, and at the knowledge metalevel which will enable effective modeling of the intelligent system itself.

3.1 The Problem Level

At the problem level in the systems management process, there are a number of abilities that an intelligent system for decision support should have. It should assist the decisionmaker in the formulation or framing of the decision situation in the sense of recognizing needs, identifying appropriate objectives by which to measure the successful resolution of an issue, and generating alternative courses of action that will resolve the needs and satisfy objectives. It should also provide support in enhancing the abilities of the decisionmaker to obtain the possible impacts on needs of the alternative courses of action,

and to understand systems behavior. This analysis capability must be associated with provision of capability to study the response or changes in the existing situation due to potential intervention as to enhance the ability of the decisionmaker to provide an interpretation of these impacts in terms of predefined and evolving objective measures. This interpretation capability will lead to evaluation of the alternatives and selection of a preferred alternative option. Associated with these must be the ability to acquire, represent, and utilize information or knowledge, and the ability to implement the chosen alternative courses of action. All of this must be accomplished with due consideration with the particular rationality perspective that is used for decisionmaking.

3.2 Cognitive Process Level

Several intelligent systems design complexities arise at the cognitive process level of systems management. These relate to the forms, frames, or perspectives associated with acquiring, integrating, and applying vast amounts of knowledge. These forms range from the systemic framework of formulation, analysis, and interpretation at the problem level that is characteristic of formal operational or holistic thought, to intuitive affect that is characteristic of concrete operational and wholistic thought. The reasoning perspectives invoked at the cognitive process level of problem resolution depend upon the task requirements, the experiential familiarity of the decisionmaker with the task, and the rationality perspectives that are used for task resolution [Sage, 1981, 1982; Linstone, 1984].

4. Knowledge Aggregation Needs in Intelligent Systems for Decision Support

Various assumptions about the nature and characteristic of the contingency structural elements of task, environment, and human problem solver familiarity with these are considered essential in the design of intelligent systems in order to enable effective and efficient organization of knowledge about specific situation domains. Due to these assumptions, not all specific domains of knowledge are suitable for building intelligent

systems for decision support. Gevarter [1984] identifies several characteristics that a domain of knowledge must satisfy at the knowledge level in order to allow an intelligent system to be built that is based only on expert knowledge:

- (1) There must be at least one human expert who is acknowledged to perform the task well.
- (2) The primary sources of the expert's abilities must be special knowledge, judgment, and experience.
- (3) The expert must be able to articulate that special knowledge, judgment, and experience and also explain the methods used to apply it to a particular task.
- (4) The task must have a well-bounded domain of application.

There will exist many situations in which these requisite conditions are not satisfied. A decision support system is a generic dual of an expert system and is intended for use in situations where at least one of the aforementioned four conditions do not apply. It appears that this will often occur. For this reason there is motivation to seek support system design incorporating features of both the expert system, whose design assumes availability of relevant expertise, and the decision support system, the design of which assumes that this expert knowledge, and its availability in a well structured format, cannot be assumed.

The representation of knowledge suggests the existence of some form of prior knowledge which enables the system to perform the function of acquisition and aggregation of the new knowledge in with the existing knowledge as to enable decision support. This much be such as to enable expansion, contraction, replacement, and residual shifts with respect to knowledge in the knowledge base. Also, notions of knowledge "quality" need to be associated with various knowledge representations and updated as new knowledge is received.

Information seeking efforts are necessarily concerned with the process by which information that is relevant to a situation is obtained from the environment. A knowledge representation system must be able to provide a description and an explanation of the situation such as to enable generation of a set of belief or knowledge, organized into a "representation," about the

situation. It should be noted that what may be a belief to one person may be knowledge to another [Abelson, 1979] and this relates to perceived information quality.

There must be also a generalization component, or inference mechanism that is equipped with some form of basic logic to enable access to the knowledge base for formation of inference and judgments. A fundamental question arises from this discussion concerning how a priori knowledge and the generalization component influence the operation of the complete intelligent system for decision support.

There are two perspectives that are relevant here concerning how individuals go about the retrieval of information, the use of information for reasoning, and the feedback process that enables acquisition of new knowledge. One perspective concerning this is that learning is performed by an elementary-to-complex process in which simple things are learned first and from this, more advanced concepts are then learned. The other perspective is based on the belief that learning starts with complex statements about the description of a situation; and that through decomposition into simpler statements a system is able to increase its understanding concerning the specific situation domain.

As a means of illustrating the two perspectives concerning the knowledge acquisition process, it is useful to compare problem solving activities and natural language understanding in some detail with respect to this two perspectives concerning learning or acquisition of new knowledge.

Most artificial intelligence systems that are in use today are based on the first perspective with respect to acquisition and aggregation of new knowledge. They use the elementary-to-complex perspective and do not typically verify, validate, or otherwise seek to determine the consistency of the resulting knowledge base. The resulting lack of control with respect to questions of validity of the resulting knowledge base is characteristic of an incomplete intelligent system. If this perspective of knowledge acquisition and aggregation is used exclusively, then some essential components of an intelligent system have been omitted, or the interaction between the support

system and the user has been modeled inadequately.

Systems that operate at the other extreme of the knowledge acquisition spectrum, such that learning proceeds in a complex-to-elementary fashion, have been difficult to implement. In reality, both modes of learning are appropriate and both are used by the human problem solver. Integration of the two approaches is clearly desirable and this is a major concern of this paper.

5 Inferential Support for Knowledge Aggregation

Inferential activities based on imprecise, incomplete, inconsistent, or otherwise imperfect knowledge is becoming more important in the design, implementation, and operation of expert systems. Inference is concerned with the generation of theories and hypotheses beyond those originally given. In planning and decisionmaking activities the information that is usually available initially is limited as to allow satisfactory performance of judgment and choice. Hence, inference is an essential activity for systems intended to aid in the learning process.

Several approaches for making inference from available information have been developed ranging from strict probabilistic Bayesian reasoning to less mathematically rigorous approaches. Analysis of systems based on these methods reveals discrepancies on the results obtained due largely to the differences in the underlying assumptions in which they are based. Quinlan [1983] contrasts several of these approaches and classifies their dissimilarities in terms of:

- (1) the way in which the uncertain information about propositions is represented,
- (2) the assumptions that form the basis for propagating information,
- (3) the control structure used for this propagation, and
- (4) the treatment of inconsistent information.

Sage and Botta [1983] also present a summary of contemporary research involving inference mechanisms in expert systems, concentrating on the extent to which these mechanisms can be Bayesian.

Most of the existing research concerning inference uses

probability theory as the standard for the representation, aggregation, and interpretation of information. However, while such theories have the advantage of modeling the uncertainties and present in human discourse, there are at least two potential difficulties. The semantic correspondence of probabilistic type expressions to natural language expressions is questionable in occasions. Also, there are other forms of information imperfections than "uncertainties" and the probabilistic representation of these other forms, while often possible, is sometimes cumbersome at best.

A large number of studies in cognitive psychology indicate that human judgments of probability values are often inconsistent with the simple axioms of probability. A comprehensive review of these efforts can be found in Sage [1981]. Often, these errors are of considerable magnitude and not just small deviations usually expected from intuitive, subjective assessments. Failure to follow the rules of probability are generally attributed to errors of application and errors of comprehension of such rules. An error of application exists if there is evidence that people know and accept a rule that they did not apply. If people do not recognize the validity of the rule they violated, it is called an error of comprehension. Since both types of errors are described in terms of violations to the rules of probability, we could as well claim that the errors are the result of a misrepresentation of human judgments about uncertainty. An error of representation refers to the semantic correspondence between the natural language expression and the symbolic representation and rules of aggregation used for inference. Errors of representation may result on a set of inconsistent hypotheses. So, an inferential inconsistency may indicate an error in representation but the contrary is not true, i.e., agreement does not necessarily reflect understanding of semantic principles. Consequently, questions arise concerning how to detect and avoid errors of representation and which framework to use in modeling uncertainty, imprecision, and potentially other forms of information imperfection as well.

Inferential activities based on logical interconnection of elements in a hierarchical net or tree are called hierarchical

inference. Hierarchical inference usually entails a series of inversion, aggregation and cascading processes to compute the likelihood of an underlying hypotheses and observable evidences based on their logical relations. Inversion involves reversing the logical relation among elements in the network in order to calculate more easily the desired relation. In a Bayesian model, the process of inversion is represented by Bayes theorem. When a datum D is perceived to have an impact on the occurrence of an event H, the relation between D and H is given by

$$P(H|D) = \frac{P(H)}{P(D)} P(D|H)$$

so the perceived effect of the likelihood of H given D is expressed in terms of the perceived effect of the likelihood of D given H. Aggregation is the task of assessing the impact of a set of data on a given hypotheses based on the immediate logical relations between the data $\{D^1, D^2, \dots\}$ and the hypotheses H. Symbolically we have $P(H|D^1, D^2, \dots) = R[P(H|D^1), P(H|D^2), \dots]$ where R is the function that aggregates the local relations $P(H|D^i)$ to form the global relation $P(H|D^1, D^2, \dots)$. Cascading is the combination of a series of immediate relations on a chain of sequential impacts to assess a global relation. For example, if a datum D is perceived to have an effect on an event E and this in turn effects $H(D \rightarrow E \rightarrow H)$, then the process of cascading consists in calculating $P(H|D)$ based on the local relations $P(H|E)$ and $P(E|D)$.

The general case of hierarchical inference involves a number of processes of inversion, aggregation and cascading. A node in the hierarchical inference net represents a finite partition of exclusive and exhaustive possible states. It may be a set of hypotheses, a set of observable or unobservable events, or more generally just data.

The impact of a given state D_i on a state A_j is given by Bayes inversion theorem as

$$P(A_j|D_i) = \frac{P(A_j) P(D_i|A_j)}{P(D_i)} \quad (1)$$

Conditioning on the states of the intermediate node B to calculate $P(D_i|A_j)$ and then inverting and cascading the result gives us

$$P(A_j|D_i) = \frac{P(A_j)}{P(D_i)} \sum_{k=1}^b P(D_i|A_j B_k) P(B_k|A_j) \quad (2)$$

In decomposition for cascading, it is usually assumed that the relation among the states of adjacent nodes is unaffected by the occurrence of states at other nodes. In this case, the likelihood of state D_i given that B_k occurred is independent of every state A_j so $P(D_i|A_j B_k) = P(D_i|B_k)$ and Eq (2) becomes for the chain of nodes $A \rightarrow C \rightarrow E$

$$P(A_j|E_1) = P(A_j) \sum_{k=1}^c \frac{P(C_k|E_1) P(C_k|A_j)}{P(C_k)} \quad (3)$$

Equation (3) is sometimes referred to as the "Modified Bayes Theorem" [Gettys and Willke, 1969] and has been used in a class of procedures called "probabilistic information processing" [Edwards et. al., 1968] to help people overcome the suboptimum behavior they show when revising probabilities of interrelated events. Use of this equation requires the assessment of large amounts of data that may be very difficult to assess intuitively in complex hierarchical inference structures. For example, the meaning of the likelihood or probability of a new state given all previous information is difficult to understand when it comprise the conjunction of a large number of states. In addition, the complexity in the processing, storing and assessment steps increases rapidly with the number of nodes in the network. This has led to the common criticism of using a formal Bayesian framework for inference [Kelly and Barclay, 1973]. Recent work, especially that by Pearl [1982ab, 1983] indicates that this criticism may not be fully justified.

An interesting, efficient scheme for the propagation of beliefs or evidence in hierarchically organized inference structures has been recently reported by Pearl [1982ab]. The scheme relies in decomposing an inference task into a series of simpler intuitive inferences linking each stage in the hierarchy to produce a global assessment. The computation of the global assessment is simplified by reformulating the general Bayesian procedure for hierarchically organized inference structures discussed here. Data can be communicated among adjacent nodes, and used to update the information at every node throughout the network. The decomposed Bayesian processing, characteristic of this scheme, allows updating to be performed by a series of local updating processes between each node and its neighbors, rather than by a central processing as in the general Bayesian framework. The likelihood of the various states of a given node depend on the entire data observed. Hence, the impact of the entire data on a given node can be decomposed in two disjoint sets of data: that obtained from the network rooted at that node and data from the network above the node. At node A let $D_d(A)$ stand for data obtained from the network rooted at A, i.e., nodes $B^1, \dots, B^L$ and nodes in the networks rooted on these; and let $D^u(A)$ be data obtained from nodes in the network above A, i.e., B and nodes above it, $A^1, \dots, A^M$ and nodes rooted on these. This decomposition prescribes how information obtained from above and below some node should be combined. A series of manipulations leads to

$$P(A_i) = a \ g(A_i) \ q(A_i) \quad (4)$$

where $g(A_i) = P(D_d(A) | A_i)$ represents the probabilistic support attributed to A_i by the nodes below it, $q(A_i) = P(A_i | D^u(A))$ represents the probabilistic support received by A_i from the nodes above it, and a is a normalization constant defined as $a = P(D_d(A) | D^u(A))^{-1}$. $P(A_i)$ is in fact a conditional probability conditioned on the existing state of knowledge.

Updating the values of g and q at every node in the light of new information allows for the calculation of the probability or likelihood of the state of every node. The calculation of g at a

node involves only data obtained from the network rooted at that node. The data obtained from the network rooted at A is equivalent to data obtained from each of the networks rooted at nodes adjacent to A. This says that g can be calculated at a node if the g 's of the nodes immediately below it and the conditional probabilities quantifying the relation between these nodes are known.

The data above A, $D^u(A)$, required to calculate $q(A_i)$, can be decomposed into two disjoint sets: $D^u(B)$ and D_d (Siblings of A). Following the same reasoning as just used, we obtain a result that enables us to compute $P(A_i)$ and $P(B_j)$ without requiring normalization. These results indicate that information to perform the local processing can be represented at each node by assessed conditional probabilities relating adjacent nodes in the hierarchy and computed values of g and $P(\cdot)$ at each node.

To initialize the inference net for propagation, we need the assessed conditional probabilities at each node. At an observational node every state is equally likely to occur in the absence of any information, hence $g(\cdot)$ is set to 1 at every observational node. From this, the value of g at every other node can be calculated. From the prior probability at the top node and the computed values of g , the probability of the states of each node can be calculated. Once the net is initialized, the occurrence of a particular state at an observational node will cause g to be updated. This information is then propagated up to update the g 's of all other nodes and then down to update the likelihood of the states of each node.

In contrast with strict Bayesian procedures, Pearl's scheme requires only the assessment of a prior probability for the node at the top of the hierarchy, that is, the last stage of the hierarchical inference structure usually representing the hypotheses being studied. The probabilities of all other stages in the structure are uniquely determined by the assessed conditional probabilities at each node, thus reducing somewhat the amount and complexity of prior information required. On the other hand, Pearl's work relies on more strict independence assumptions in order to obtain computationally tractable results, and also requires prior knowledge about the distribution of the

underlying hypothesis being studied.

One of the major criticisms of this, and similar Bayesian approaches, is the need to identify point values about the probability of events. Usually, a point value assessment of the probability of an event is an overstatement about our actual knowledge of the likelihood of occurrence of that particular event. In response to the need of representing imprecision of Bayesian probability values, Dempster [1967] utilized the concept of lower and upper probabilities to deal with the subjective imprecision of uncertainty measures. Shafer [1976, 1981] presents a comprehensive exposition of this novel idea as well as extensions to the theory of inference based on the concept of upper and lower probabilities. The basic idea of this concept is that instead of representing the probability of an event A by a point value $P(A)$, it may be bounded by a subinterval of $[0,1]$. That is, the exact probability $P(A)$ may be unknown but bounded. This kind of representation has solid grounds in the Dempster-Shafer theory of basic probability and for that reason has received considerable attention recently.

Of particular interest in this research is the work of Toulmin et. al. [1979] in that they have constructed an explicit structured model of logical reasoning that is suited for analytical inquiry and computer implementation. The model is sufficiently general that it can be used to represent logical reasoning in a number of application areas.

Starting from the assumption that whenever we make a claim there must be some grounds in which to base our conclusion, Toulmin states that our thoughts are generally directed from the GROUNDS to the CLAIM. The GROUNDS and the CLAIM are statements that express fact and values. As a means of stating observed patterns of stating a claim, there must be a reason that can be identified to connect the GROUNDS and the CLAIM. This connection is called the WARRANT, and it is the WARRANT that gives to the GROUNDS-CLAIM connection its logical validity.

We say that the GROUNDS support the CLAIM on the basis of the existence of a WARRANT that explain the connection between the grounds and the claim. It is easy to relate the structure of this basic elements with the process of inference, whether

statistical, deductive, or inductive. The WARRANTS are the set of rules of inference, and the GROUNDS and CLAIM are the set of well defined propositions or hypotheses. It will be only the sequence and procedures, that are used to come up with the three basic elements and their structure in a logical fashion, what will determine the type of inference that is used.

Sometimes, in the course of reasoning about an issue, it is not enough that the WARRANT will be the absolute reason to believe the CLAIM on the basis of the GROUNDS. For that, Toulmin allows for further BACKING which, in his representation, supports the WARRANT. It is the BACKING that provides for the reliability, in terms of truth, associated with the use of the WARRANT. The relationship here is analogous to the way in which the GROUNDS support the CLAIM. An argument will be valid and will give the CLAIM solid support only if the WARRANT is relied upon and is relevant to the particular case under examination. The concept of logical validity seems to imply that we can only make a CLAIM when both the WARRANT and the GROUNDS are certain. However, imprecision and uncertainty in the form of exceptions to the rules or low degree of certainty in both the GROUNDS and the WARRANT does not prevent us on occasions from making a "hedge" or a vague CLAIM. Very commonly, we must arrive at conclusions on the basis of something less than perfect evidence; and we put those claims forward not with absolute and irrefutable truth but rather with some doubt or degree of speculation.

To allow for these cases, Toulmin adds MODAL QUALIFIERS and POSSIBLE REBUTTALS to his framework for logical reasoning. MODAL QUALIFIERS refer to the strength or weakness with which a claim is made. In essence every argument has a certain modality. Its place in the structure presented so far must reflect the generality of the WARRANT in connecting the GROUNDS to the CLAIM, and also with condition of validity of the set of facts as GROUNDS. POSSIBLE REBUTTALS, on the other hand, are exceptions to the rules. Although MODAL QUALIFIERS serve the purpose of weakening or strengthening the validity of a CLAIM, there may be still conditions that invalidate either the GROUNDS or the WARRANTS, and this will result in deactivating the link between the CLAIM and the GROUNDS. These cases are represented by the

POSSIBLE REBUTTALS.

The resulting structure of logical reasoning provides a very useful framework for the study of human information processing activities. The order in which the six elements of logical reasoning has been presented serve only the purpose of illustrating their function and interdependence in the structure of an argument about a specific issue. It does not represent any normative pattern of argument formation. In fact, due to the dynamic nature of human reasoning, the concept formation and framing that results in a particular structure may occur in different ways. The six element model of logical reasoning is shown in Figure 1.




Figure 1. The six element model of logical reasoning

The effects of various forms of inquiry upon issues of representation and detection of judgmental errors in human information processing have been investigated [Lagomasino and Sage 1985ab, 1986] using this structure of rational argument. The frameworks for Bayesian inference just discussed requires probability values as primary inputs. Since most events of interest are unique or little is known about their relative frequencies of occurrence, the assessments of probability values usually require human judgment. Substantial psychological research has shown that people are unable to elicit probability values consistent with the rules of probabilities and that they are suboptimum in revising probability assessment when new

information is obtained.

Tversky and Kahneman [1982] have shown that dominance of causal over diagnostic information exists in assessing conditional probabilities. They concluded, in a series of experiments, that if some information have both causal and diagnostic implications then people, instead of weighting the causal and diagnostic impacts of the evidence, they apparently assess conditional probabilities primarily in terms of the direct causal effect of the impacts. Thus, if A is perceived to be the cause of B, people will give higher probabilities to $P(B|A)$ than to $P(A|B)$. Burns and Pearl [1981] conducted a study to test the validity of judgments made by these two forms of reasoning. Thus, they investigated whether causal or diagnostic judgment is a more natural way of encoding knowledge about every day experiences. Their results demonstrated that neither one was found to be more accurate than the other. In a similar study, Moskowitz and Sarin [1983] reported that individuals found it easier and showed more confidence in assessing $P(A|B)$ if B is causal to A.

This apparent contradiction of results may be explained by the differences in the contingency task structure within which the experiments were performed, and that suggest that the choice of which form of inference to invoke depends more on the level of familiarity of the observer with the task at hand and the cognitive style which determines the way in which the knowledge was originally perceived. Most structuring procedures for decision making rely on the "divide and conquer" approach under the assumption that judgment is improved when a complex, ill-defined problem is decomposed, analyzed and solved by a set of smaller, well-defined problems. The findings of these studies, aside from having implications on the validity of the "divide and conquer" approach, implies that the form of representation of judgments used should correspond with the meaning of the judgments assessed.

Falk and Bar-Hillel, [1979] have recognized the importance of distinguishing between probabilistic and logical support. Probabilistic support refers to the increased in likelihood of the occurrence of an event A given that another event B has occurred. That is, A supports B if $P(B|A) > P(B)$. Logical support

exhibits the relation of implication between two premises, denoted $A \rightarrow B$, that fails to hold only if the first is true and the second is false. Logical support is transitive; if $A \rightarrow B$ and $B \rightarrow C$, then $A \rightarrow C$. When $A \rightarrow B$ it is also true that $B \rightarrow A$. With these definitions, the distinctions between probabilistic and logical support should be apparent. Logical support (ATB) does not imply conditions similar to those that follow from probabilistic support. $A \rightarrow B$ is logically equivalent to $B \rightarrow A$, but does not say anything about the truth of $A \rightarrow B$ or $B \rightarrow A$. Likewise, probabilistic support is not transitive and logical support must be transitive. A major point in this distinction that arises here is when to apply these two methods of representation in inferential activities. We are concerned with this issue because, as previous research has shown, the method used to represent human judgments may strongly influence the validity and consistency of these judgments.

It is especially important here to note that we can ascribe uncertainty to the truth of logical support by referring to the probability by which A implies B , $P(A \rightarrow B)$. The importance of this will concern us next in the development of a general framework for inference. Much of our research has been concerned with developing a framework of inference suitable for assessing and structuring complex problems that derives from the logic of reasoning of Toulmin and the calculus of probabilities to make inferences on the likelihood of the events or premises that comprise the inference structure. Assessments by the decision maker in the form of logical support relations among the events are used to structure the problem, and assessments in the form of set inclusion inequalities among the events and their relations are combined using the probability calculus to draw inferences.

The framework and process for inference support developed here is applicable to a general class of networks of interrelated propositions. Specifically, it can be applied to finite connected networks where the number of propositions is finite and where every pair of distinct propositions is joined by at least one chain of relations. We have developed a procedure that describes the information processing elements involved in the structuring and analysis of an issue within this framework. The

information processing functions associated with the use of the framework for inference described here involves four steps:

- 1) initial problem framing,
- 2) hypothesis generation,,
- 3) parameter value assessment, and
- 4) hypothesis evaluation and situation assessment.

The intent of the first step is to capture those elements and relations that constitute an issue, and to represent them in a form that is suitable for inference. The inference network developed here are not necessarily hierarchical as contrasted with the case of Bayesian inference. We are able to deal with structures that correspond to a very general type of inference network. Nodes in the network represent the propositions of interest in the particular issue. Inferential links between propositions are defined in terms of the set of consistency relational equations, including the set of consistency relations and any other assessed relations between the propositions at each node. The probability value of the propositions at each node is underconstrained and acquisition of information about the relation between the nodes is the primary means of further constraining the probability values of the propositions at each node.

Given the assessed initial problem frame, the task of hypothesis generation involves the generation of reasonable hypotheses that are based on situational perception and information needed for the task at hand. In most cases this involves the specification of alternative hypotheses at each node. Ideally, the set of hypotheses under consideration at each node should be mutually exclusive and exhaustive. This task also involves the selection of the basic premises and possible rebuttals relevant for each inferential link.

The parameter value identification step provides for the continual assessment of the parameters of the inference model. This includes the assessment of probability values of the propositions at each node as well as the probability values representing the uncertain logical relations at each inferential link. These assessments can be related and represented imprecisely in the form of bounded intervals and/or linear

inequalities on the set of parameters. Achieving the task of parameter value assessment with minimal imprecision will depend strongly on the quality of the information available and the person's perception and familiarity with the task at hand.

The hypothesis evaluation and situation assessment step involves probability categorization, over a set of alternative hypotheses, of the probable situations as captured by the information that is provided to the inference model. Given the set of consistency relational equations for each link in the inference network, we can calculate the probability values for the propositions at each node. These probability values will usually be stated in the form of bounded intervals and linear inequalities. If more precise information is required then further assessments about the relations and propositions in the network must be performed. This suggests the generation of alternative hypotheses and the assessment of more precise parameter values. The information processing tasks required in the use of the framework for inference based on logical reasoning describes an iterative process suitable for situations where knowledge about it is ill defined or imperfectly described.

The objective of this portion of our research has been to investigate various approaches for inference based on imprecise knowledge, and to advance the state of research in the area of representation of natural language expressions about uncertainty and imprecision. We have investigated a new approach for inference based on logical support relations that differs considerably from Bayesian approaches which rely on probabilistic support relations. This new approach has the interesting feature of being computationally simple, capable of working in a general class of inference networks, not relying on idealistic independence assumptions, and not having to make a clear distinction between hypothetical and evidential type of information.

6. Effectiveness Issues and Conclusions

In the context of devices to aid in human information processing activities, there exist two design philosophies concerning the proper relationship between the input component

and the generalizing or inference mechanism. One advocates the view that these components should be considered as operating separately; with the input component in charge of knowledge acquisition and representation considered as if it was independent of the generalizing sector which is in charge of aggregation or inference. This approach offers simplicity in system design, at the expense of effectiveness. The mode of representation of information will influence the success or failure of the inquirer in arriving at a solution. This deficiency seems to be characteristic of many current information systems. They are passive and it is up to the user to recognize an information need and then seek out the required information.

The other approach considers that the two sectors are essentially nonseparable and that each supports and enhances the functioning of the other; but this in turns complicates the system design, perhaps by a considerable amount. In this approach, the internal interactions of the input and generalizing sector are capable of generating user-system interaction. There are various ways in which a system can initiate a dialog with a user. These include identification of:

- (1) "gaps" in the knowledge base that prevent the system from making inferences, or from adequately summarizing the information in a sector of the knowledge base;
- (2) an inconsistent set of information followed by detection of the inability of the system to resolve it; or
- (3) sufficiently imprecise information that makes it not possible to suggest an appropriate course of action; and
- (4) inability of the system to satisfy the desired goals of the user.

Identification of these potential deficiencies, and use of prompts based on them for purpose of computer-control dialog are needs in intelligent system design for decision support. To be fully responsive to the needs of users research concerning this should be integrated with research involving inference structures.

In this paper and in the associated research, we have been

especially concerned with situations in which there are a variety of perspectives and experiential familiarity with a particular judgmental task. It is argued that this requires attention to combining the skill based and rule based features of an expert system with the formal reasoning based knowledge developed through use of a decision support system. This could be especially fruitful in accomplish meta-level model based management. A structured inference procedure to accomplish this type of control, based on the use of imperfect information as well as probabilistic and logical support, was described.

6. References

Abelson, R. P., "Differences Between Belief and Knowledge Systems," Cognitive Science, Vol. 3, 1979, pp. 355-366.

Barr, A., and Feigenbaum, E. A., (EDS) The Handbook of Artificial Intelligence, Vol. 1 & 2, William Kaufman, Inc., Los Altos, CA, 1981, 1982.

Burns, M. and Pearl, J., "Causal and Diagnostic Inferences: A Comparison of Validity," Organizational Behavior and Human Performance, Vol. 28, 1981, pp. 379-394.

Cohen, P. R., and Feigenbaum, E. A., (EDS) The Handbook of Artificial Intelligence, Vol. 3, William Kaufman, Inc., Los Altos, CA, 1982.

Date, C. J., An Introduction to Database Systems, Addison-Wesley Publishing Co. Reading, MA., 1981.

Date, C. J., Database: A Primer, Addison-Wesley Publishing Co. Reading, MA., 1983.

Dempster, A. P., "Upper and Lower Probabilities Induced by a Multivalued Mapping," Annals of Mathematical Statistics, Vol. 38, 1967, pp. 325-339.

Edwards, W., Phillips, L. D., Hays, W. 1. and Goodman, B.C., "Probabilistic Information Processing Systems: Design and Evaluation," IEEE Transactions on System Science and Cybernetics, Vol. SMC-4, 1968, pp. 248-265.

Falk, R. and Bar-Hillel, M., "Probabilistic Dependence between Events," Hebrew University of Jerusalem, October 1979.

Gettys, C. and Willke, T. A., "The Application of Bayes' Theorem When the True Data State is Uncertain," Organizational Behavior and Human Performance, Vol. 4, 1969, pp. 125-141.

Gevarter, W. B. Artificial Intelligence, Expert Systems, Computer Vision and Natural Language Processing, Noyes Publications, 1984.

Ishizuka, M. "Inference Methods Based on Extended Dempster-Shafer Theory for Problems with Uncertainty/Fuzziness," New Generation Computing, Vol. 1, 1983, pp. 159-168.

Kahneman, D. and Tversky, A., "Variants of Uncertainty," Cognition, Vol. 11, No. 2, March 1982, pp. 143-157.

Kahneman, D., Slovic, P., and Tversky, A. (Editors), Judgments under Uncertainty: Heuristics and Biases, Cambridge University Press, New York, 1982.

Kelly, C. W. and Barclay, S., "A General Bayesian Model for Hierarchical Inference," Organizational Behavior and Human Performance, Vol. 10, 1973, pp. 388-403.

Lagomasino, A. and Sage, A. P., "Imprecise Knowledge Representation in Inferential Activities," in M. M. Gupta (ED) Approximate Reasoning in Expert Systems, North Holland Elsevier, 1986.

Lagomasino, A. and Sage, A. P., "An Interactive Inquiry System," Large Scale Systems, Vol. 9, 1985a.

Lagomasino, A., and Sage, A. P., "Representation and Interpretation of Information for Decision Support with Imperfect Knowledge," Large Scale Systems, Vol. 9, 1985b.

Linstone, H. A., Multiple Perspectives for Decision Making: Bridging the Gap Between Analysis and Action, Elsevier North-Holland, New York, 1984.

Moskowitz, H. and Sarin, R. K., "Improving the Consistency of Conditional Probability Assessments for Forecasting and Decision Making," Management Science, Vol. 29, No. 6, June 1983, pp. 735-749.

Pearl, J., "Distributed Bayesian Processing for Belief Maintenance in Hierarchical Inference Systems," UCLA Cognitive Systems Lab. UCLA-ENG-CSL-82-11, January, 1982.

Pearl, J., "Reverend Bayes on Inference Engines: A Distributed Hierarchical Approach," presented at the AAAI National Conference on Artificial Intelligence, August 18-20, 1982, Pittsburgh, Pennsylvania.

Pearl, J., Heuristics: Partially Informed Strategies for Computer Problem Solving, Addison-Wesley, Reading, Mass., 1983.

Quinlan, J. R., "Inferno: A Cautious Approach to Uncertain Inference," The Computer Journal, Vol. 26, No. 3, 1983, pp. 255-266.

Sage, A. P., "Behavioral and Organizational Considerations in the Design of Information Systems and Processes for Planning and Decision Support," IEEE Transaction on Systems, Man, and Cybernetics, Vol. SMC-11, No. 9, September, 1981, pp. 640-678.

Sage, A. P. and Botta, R., "On Human Information Processing and Its Enhancement Using Knowledge-Based Systems," Large Scale Systems, Vol. 5, 1983, pp. 208-223.

Sage, A. P. and Lagomasino, A., "Knowledge Representation and Man Machine Dialogue," a Chapter in Advances in Man-Machine Systems Research, William B. Rouse, (ED), JAI Press, 1984, pp. 223-260.

Sage, A. P., and White, C. C., "ARIADNE: A Knowledge Based Interactive System for Planning and Decision Support," IEEE Transactions on Systems, Man and Cybernetics, Vol. 14, No. 1, January 1984, pp. 35-47.

Shafer, G., A Mathematical Theory of Evidence, Princeton University Press, Princeton, N.J., 1976.

Shafer, G., "Constructive Probability," Synthese, Vol. 48, 1981, pp. 1-60.

Sprague, R. H. Jr., and Carlson, E. D., Building Effective Decision Support Systems, Prentice Hall, Inc., Englewood Cliffs, NJ, 1982.

Toulmin, S., Rieke, R. and Janik, A. An Introduction to Reasoning, McMillan Publishing Co., 1979.

Tversky, A. and Kahneman, D., "Causal Schemas in Judgments under Uncertainty," in Kahneman, D., Slovic, P., and Tversky, A. (EDS), Judgments under Uncertainty: Heuristics and Biases, Cambridge University Press, New York, 1981.