

Technical Report 576

12

AD-A135519

KNOWLEDGE REQUIREMENTS AND MANAGEMENT IN EXPERT DECISION SUPPORT SYSTEMS FOR (MILITARY) SITUATION ASSESSMENT

Moshe Ben-Bassat and Amos Freedy

The Israel Institute of Business Research

BASIC RESEARCH

DTIC
ELECTE
DEC 8 1983
S B



U. S. Army

Research Institute for the Behavioral and Social Sciences


August 1983

Approved for public release; distribution unlimited.

DTIC FILE COPY

83 12 08 012

DISCLAIMER NOTICE



THIS DOCUMENT IS BEST QUALITY AVAILABLE. THE COPY FURNISHED TO DTIC CONTAINED A SIGNIFICANT NUMBER OF PAGES WHICH DO NOT REPRODUCE LEGIBLY.

**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**

**L. NEALE COSBY
Colonel, IN
Commander**

**Research accomplished under contract
for the Department of the Army**

The Israel Institute of Business Research

Technical review by

**Robert M. Sasmor
Joseph M. Piljay**

NOTICES

DISTRIBUTION: Primary distribution of this report has been made by ARI. Please address correspondence concerning distribution of reports to: U.S. Army Research Institute for the Behavioral and Social Sciences, ATTN: PERI-TST, 5001 Eisenhower Avenue, Alexandria, Virginia 22333.

FINAL DISPOSITION: This report may be destroyed when it is no longer needed. Please do not return it to the U.S. Army Research Institute for the Behavioral and Social Sciences.

NOTE: The findings in this report are not to be construed as an official Department of the Army position, unless so designated by other authorized documents.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER Technical Report 576	2. GOVT ACCESSION NO. ADA135519	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Knowledge Requirements and Management in Expert Decision Support Systems for (Military) Situation Assessment	5. TYPE OF REPORT & PERIOD COVERED INTERIM REPORT	
	6. PERFORMING ORG. REPORT NUMBER 46/81	
7. AUTHOR(s) Dr. Moshe Ben-Bassat and Amos Freedy	8. CONTRACT OR GRANT NUMBER(s) DAJA37-81-C-0065	
	9. PERFORMING ORGANIZATION NAME AND ADDRESS The Israel Institute of Business Research Tel Aviv University; University Campus Ramat-Aviv, PO Box 39010, Tel-Aviv 69978	
11. CONTROLLING OFFICE NAME AND ADDRESS USARDCG (UK) Box 65 FPO New York, NY 09510	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 2Q161102B-74D	
	12. REPORT DATE August 1983	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) US Army Research Institute for the Behavioral and Social Sciences 5001 Eisenhower Avenue Alexandria, VA 22333	13. NUMBER OF PAGES 46	
	15. SECURITY CLASS. (of this report) UNCLASSIFIED	
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 Technically monitored by Dr. Michael Kaplan		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Decision support systems, expert systems, knowledge elicitation and representation, situation assessment, classification models.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Situation assessment tasks, e.g. medical diagnosis, battlefield reading, corporation assessment for merger or acquisition purposes, are formulated as a general family of problem solving tasks. We characterize the generic nature of this family as a multiperspective multimembership hierarchical pattern recognition problem, identify the types of decision problems involved in the situation assessment process, and propose a unified approach for the development of situation assessment decision support systems (DSS). The focus		

→ all

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

1 SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

crit

is on knowledge representation and elicitation, although issues related to inference mechanisms, system structure and expert-machine-user interface are also discussed.

The presentation is accompanied by examples from military situation assessment. However, comparable examples from medical and business applications are also cited. Many of the ideas presented here have already been implemented in the MEDAS system; a medical DSS for emergency and critical care medicine.

UNCLASSIFIED

Technical Report 576

**KNOWLEDGE REQUIREMENTS AND MANAGEMENT
IN EXPERT DECISION SUPPORT SYSTEMS
FOR (MILITARY) SITUATION ASSESSMENT**

Moshe Ben-Bassat and Amos Freedy

The Israel Institute of Business Research

Contract monitored by Michael Kaplan

**Approved as technically adequate
and submitted for publication by
Robert Seemor, Director
BASIC RESEARCH**

**U.S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES
8001 Eisenhower Avenue, Alexandria, Virginia 22333**

**Office, Deputy Chief of Staff for Personnel
Department of the Army**

August 1983

**Army Project Number
2Q161102B74D**

Basic Research

Approved for public release; distribution unlimited.

ARI Research Reports and Technical Reports are intended for sponsors of R&D tasks and for other research and military agencies. Any findings ready for implementation at the time of publication are presented in the last part of the Brief. Upon completion of a major phase of the task, formal recommendations for official action normally are conveyed to appropriate military agencies by briefing or Disposition Form.

BRIEF

Situation assessment tasks, e.g. medical diagnosis, battlefield reading, corporation assessment for merger or acquisition purposes, are formulated as a general family of problem solving tasks. We characterize the generic nature of this family as a multiperspective multimembership hierarchical pattern recognition problem, identify the types of decision problems involved in the situation assessment process, and propose a unified approach for the development of situation assessment decision support systems (DSS). The focus is on knowledge representation and elicitation, although issues related to inference mechanisms, system structure and expert-machine-user interface are also discussed. Two types of knowledge are distinguished; global knowledge and local knowledge. Global knowledge is required to determine directions on which to focus attention, while local knowledge is required for assessing the validity of a specific alternative based on a given set of findings. Global knowledge is represented as a network of relevancy pointers between alternatives and features. Attached to the links of this network are weights by which the strength of relevancy is evaluated and global directions (hypotheses) for situation analysis are determined. For local knowledge, it seems that in most practical problems multiple representation techniques would be required to characterize adequately the alternatives by means of their relevant features.

The presentation is accompanied by examples from military situation assessment. However, comparable examples from medical and business applications are also cited. In fact, many of the ideas presented here have already been implemented in the MEDAS system; a medical DSS for emergency and critical care medicine.

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/ _____	
Availability Codes	
Dist	Avail and/or Special
A-1	



v

**KNOWLEDGE REQUIREMENTS AND MANAGEMENT IN EXPERT DECISION SUPPORT SYSTEMS
FOR (MILITARY) SITUATION ASSESSMENT**

CONTENTS

	Page
INTRODUCTION	1
SITUATION ASSESSMENT - A TECHNICAL FORMULATION	4
TAXONOMY OF DECISION PROBLEMS INVOLVED IN SITUATION ASSESSMENT	7
BASIC CHARACTERISTICS OF AN EXPERT DECISION SUPPORT SYSTEM	11
REQUIREMENTS ANALYSIS FOR KNOWLEDGE REPRESENTATION AND ELICITATION	13
AN APPROACH FOR KNOWLEDGE REPRESENTATION AND ELICITATION	16
SITUATION ASSESSMENT INFERENCE ALGORITHMS	27
SYSTEM STRUCTURE AND HUMAN-MACHINE COMMUNICATION	28
SUMMARY	32
REFERENCES	35

TABLES

Table 1. A Pattern for Attack Intention	22
2. Interval Estimates for Quantifying Conditional Probabilities	23

FIGURES

Figure 1. Multiperspective hierarchical structure of the situation assessment problem	5
2. System block diagram	29

1. Introduction

Situation assessment, i.e. knowing where we are, constitutes a fundamental problem in many decision making processes. In business management we face this problem when determining the status of an organization from one or more perspectives, such as financial stability, competitors threat, productivity performance and marketing performance. Medical diagnosis is a similar problem in which the patient situation is assessed with respect to several sets of potential disorders in the various body systems. Another example is military situation assessment, a problem which is also known as battlefield reading.

All of these problems share common characteristics in that the decision maker starts with some uncertainty with regard to the true situation, and then looks for additional information which may reduce this uncertainty. Following a cyclic process new information is obtained and integrated into the existing information, the situation is reassessed and if final assessment cannot be made, further information is requested. The process ends when the decision maker decides that he knows enough about the situation and can make up his mind, or that no additional sources of information can contribute significantly (compared to their cost) to remove the uncertainty which still remains, or temporal considerations force him to terminate information acquisition and assess the situation as best he can.

Computer-based decision support systems (DSS) for situation assessment tasks have been widely proposed in the past for various applications, including medicine [1], weather forecasting [2] and military [3]. The classical approach to these systems is based on simple classification models. By this approach the object whose situation is to be assessed (e.g. patient, battlefield, corporation) is represented by a vector of features which characterize this object with

respect to a given set of classes (e.g. diseases, enemy intentions, financial status). The decision models typically used to classify a given object include the classical Bayesian model [4], [5], template matching and other statistical pattern recognition models such as nearest neighbor, and discriminant analysis [6].

Such an approach is useful for problems with fairly limited scope and complexity. For instance, the above models are not directly applicable to situation assessment problems where:

1. The classes are not mutually exclusive and exhaustive (e.g. several diseases may simultaneously exist in a given body system).
2. Complete assessment of the situation involves several interrelated classification schemes (e.g. patient's situation in several body systems has to be assessed).
3. The significance of features for recognizing the classes is not straightforward but rather via a chain of inference relationships, e.g. hierarchical.
4. The characterization of the classes involve structural and temporal relationships among the features.

In addition, the representation of the problem as a simple classification problem provides the system with very limited "understanding" of the problem structure. This may be sufficient to generate meaningful interpretation of the findings, including probability estimates regarding the true situation. However, it substantially limits the system in its ability to handle subtle cases and to discuss with the user the reasoning behind its interpretation.

Decision support systems (also known as consultation systems or expert systems) where these deficiencies were partially corrected include the MYCIN system for bacteria identification and treatment [7], the INTERNIST system (now named CADUCEUS) for internal medicine [8], the PROSPECTOR system for

mineral exploration [9] and the MEDAS system for decision support in emergency and critical care medicine [10]. Additional approaches are reviewed by Michie [11], Kulikowski [12], Shortliffe et al [13] and Gomez and Chandrasekaran [14].

As can be seen from the above references, most of the research toward situation assessment decision support systems was oriented to specific application areas, and particularly to medical diagnosis. Few attempts were made to generalize DSS systems, among them the use of the MYCIN rule-based engine to other problem areas [15] such as fault diagnosis of cars [16]. These generalizations were typically motivated by the technique, i.e. the identification of additional decision problems for which the MYCIN engine may also be used. In this paper a generalization of DSS systems is attempted from the decision process point of view. We characterize the generic nature of a general situation assessment process, identify the types of decision problems involved in this process, and propose a unified approach for the development of situation assessment support systems. The focus is on knowledge representation and elicitation, although issues related to inference mechanisms, system structure and expert-machine-user interface are also discussed. The presentation is accompanied by examples from military situation assessment. However, comparable examples from medical and business applications are also cited. In fact, many of the ideas presented here have already been implemented in the MEDAS system.

In Section 2 a detailed analysis of the situation assessment process is presented. From this analysis we derive in Section 3 the types of decision problems involved in the situation assessment process. Section 4 discusses the basic characteristics and elements of expert decision support systems. The requirements for knowledge representation and elicitation are discussed in Section 5, while Section 6 proposes an approach for this task. Inference algorithms are discussed in Section 7 and Section 8 describes system structure. Section 9 concludes with summary remarks.

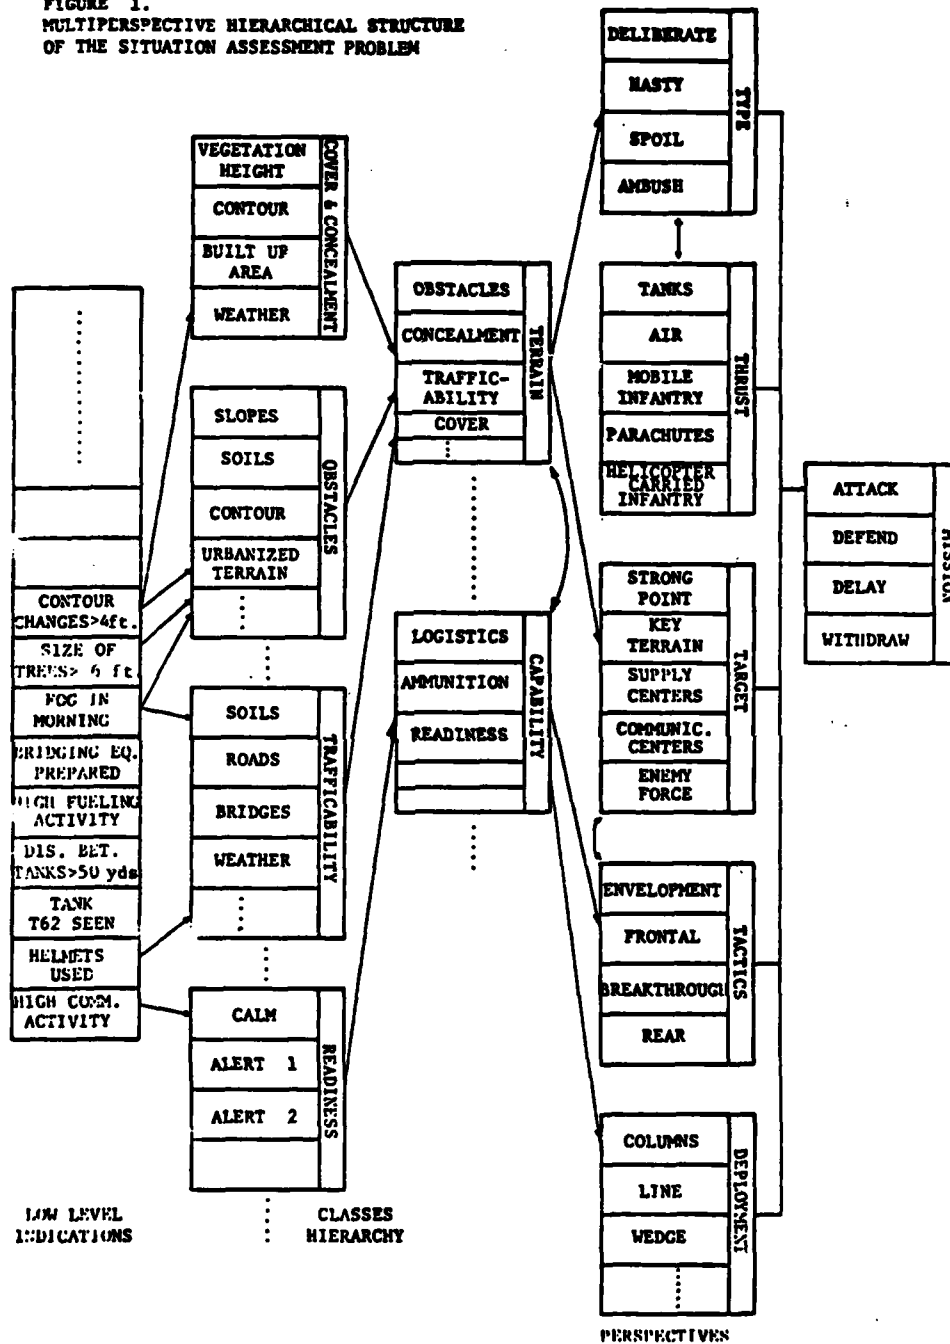
2. Situation Assessment - A Technical Formulation

Situation assessment may be considered as a multi-perspective multi-membership hierarchical pattern recognition problem. Figure 1 explains these terms. It shows a possible representation of the relationships between military indicators which may be used by the situation assessment process. The recognition process is "multi-perspective" in the sense that the overall picture of the situation is constructed from elements recognized in various perspectives of the battlefield. For instance, the various perspectives for analyzing an attack include TYPE, THRUST, TARGET, TACTICS, DEPLOYMENT, etc. In each of these interrelated perspectives, the situation may be classified within one or more of the alternatives (classes) associated with that perspective. For example, in Figure 1, an enemy attack can be one of the following TYPES: DELIBERATE, HASTY, SPOILING, or an AMBUSH. Similarly, there are several alternatives for THRUST, TARGET, TACTICS, etc.

The term "multi-membership" refers to the possibility that within a given perspective several alternatives may co-exist simultaneously. For instance, within the THRUST perspective there is no reason to assume a priori that the enemy attack will consist of TANKS only or PARACHUTES only. Any combination of the possible alternatives: TANKS, AIR, MOBIL INFANTRY, PARACHUTES, HELICOPTER CARRIED INFANTRY; may simultaneously be true.

The recognition process is "hierarchical" in the sense that low-level indications are used as the building blocks of higher level indications. For instance, information regarding the presence of trees, their height and density are features that contribute to determine COVER and CONCEALMENT. Boulder size and soil type contribute to determine tank TRAFFICABILITY. Together they contribute to TERRAIN analysis. The results of TERRAIN analysis and other factors such as CAPABILITY contribute, in turn, to the determination of what TACTICS the enemy may choose, his DEPLOYMENT technique, and even influence the choice of a TARGET.

FIGURE 1.
MULTIPERSPECTIVE HIERARCHICAL STRUCTURE
OF THE SITUATION ASSESSMENT PROBLEM



The recognition process is mostly directed bottom-up. Occasionally, however, correlations between indications at the same level may provide horizontal evidence as well. For instance, recognizing the tactics of an attack may suggest evidence regarding the target of the attack and vice versa. This is indicated in Figure 1 by horizontal links. Moreover, it is also possible that certain indications provide evidence for other indications at lower levels. For instance, if indications associated with "presence of a reconnaissance battalion" clearly indicate that a division level attack is anticipated, this may increase the likelihood that an observed and previously hard to interpret column of tanks is a part of the division attack, and even its target becomes easy to determine.

3. Taxonomy of Decision Problems in Situation Assessment

Figure 2 illustrates the cycle through which the situation assessment process iterates. Each step in this cycle represents one type of a decision problem each of which may require different skills and techniques.

(0) Initial Findings Accumulation

The cycle starts with the presentation of an initial set of specific facts about the situation. These facts may have been observed in the field or they may have been passed to the decision maker, e.g. a G2 officer, through the command channels. They may have come from higher echelons, from parallel units or from subordinate units. They also include indications or responses to information requests that have been placed previously by the G2 and collected by the various information collecting agencies at his disposal.

From thereon the process may be decomposed into the following steps:

(1) Hypotheses Generation and Evaluation

The findings recently obtained are integrated into the existing evidence (which in the first iteration is the apriori information only) and trigger a moving chain of deductions pointing at several alternative classes in several perspectives of the battlefield. The uncertainty regarding the truth of these classes is updated, and as a result some classes may be verified (beyond some threshold of confidence), others may be refuted (below some reasonable threshold of confidence) and still others may remain uncertain, though still feasible. At this point an attempt is made to see if the entire puzzle is clear, i.e. if the existing evidence explains the situation in each perspective of the battlefield and a global interpretation of the situation may be drawn. Those aspects of the battlefield which remain unclear serve as the basis for deriving hypotheses to be worked up in the subsequent stages. The generation of a rich set of plausible hypotheses is the hallmark of a good situation assessor.

(2) Goal(s) Setting

Occasionally - particularly in early stages - too many hypotheses may be triggered by the existing evidence, and not all of them may be simultaneously explored. In such a case goal(s) need to be set on which attention will be focused in the next immediate stages. These may include, for instance, verification/elimination of a specific hypothesis, or differentiation between a group of competing hypotheses. Factors which affect goal determination include the severity and urgency of the candidate alternatives (e.g. enemy attack is expected within 24 hours) their present level of uncertainty and their initial apriori incidence [35].

(3) Information Sources Evaluation and Selection

Once a goal(s) is set, the information sources which may offer the findings by which this goal may be achieved, need to be identified and evaluated. Such an evaluation is based, on one hand, on the potency (information content and reliability) of these information sources to achieve the determined goal, and, on the other hand, on the cost of utilizing them. This cost reflects not only financial, technical and logistic investments, but also the risk involved in getting the information. The information source(s) with the greatest expected contribution to the specified goal(s) compared to its cost is then invoked, e.g. a reconnaissance aircraft. Frequently, a battery of information sources may be utilized simultaneously to permit deeper exploration of a given hypothesis or concurrent exploration of several hypotheses.

(4) Findings Sorting by Goals and Hypotheses

As new findings come in, either as a result of the decision maker's request or "voluntarily", they should be sorted with regard to the entire battlefield structure, the triggered hypotheses and, on the highest priority, with regard to the current goals. Findings should not be ignored just because they do

not contribute to the current goal(s) or to the previously activated hypotheses. It is such a lateral thinking that may open new ideas leading to the generation of new hypotheses which may eventually turn out to include the correct ones. Goals need to be set in order to direct effectively the information acquisition path. However, once a finding is observed, its significance should be analyzed with respect to all of its relevant classes.

(5) Evidence Integration

Once all of the relevancy links of the new findings are identified, they are integrated with the existing findings and not just added to them. Recognizing dependencies between new and existing findings may prevent artificial compounding of redundant information. It may also suggest synergy, i.e. the evidence suggested by the group of findings is greater than the sum of the individual findings' evidence. At this stage we may also try to restructure the grouping of findings in an attempt to discover new possible interpretations. The new integrated evidence modifies the uncertainty of existing hypotheses and may suggest new hypotheses concerning the true situation. This completes the cycle and brings us back to stage (1) unless the termination test is positive.

(6) Termination

The situation assessment cycle is interrupted or fully terminated under one of the following conditions:

- (a) A decision may be reached with regard to the true situation in each aspect of the battlefield, all of the (suspicious) findings are explained by this interpretation, and no additional hypotheses are sufficiently triggered to justify further exploration.
- (b) Several triggered hypotheses have not yet been settled, however, the cost of removing the remaining uncertainty is relatively high compared to the expected information gain and the impact on the battle plan (or treatment plan if medical diagnosis is the case).

- (c) New developments (e.g. sudden enemy attack) forces the decision maker to terminate information acquisition and assess the situation as best as he can with the existing evidence.

(7) Summary Composition

The situation assessment process culminates in the composition of the individual decisions made for separate battlefield aspects into one complete and coherent picture that leads to tactical planning. The end result is the Intelligence Estimate document which is currently produced manually by the intelligence officer.

4. BASIC CHARACTERISTICS OF AN EXPERT DECISION SUPPORT SYSTEM

Competent situation assessment requires lateral as well as vertical thinking. Lateral thinking [17] is required in order to generate "all" of the possible interpretations of a given set of findings, or to identify all of the potential sources of information for resolving a given hypothesis. Vertical thinking is required in order to explore thoroughly the significance of a given set of findings, or to assess the potency of a given information source. Being competent in each of the tasks involved in the situation assessment process, requires quite a unique set of skills, tools, experience and knowledge of the problem domain.

The ideal decision support system is capable of offering assistance in each of the tasks involved. It operates as an assistant who serves all the user's requests for information retrieval and computation, and also as an advisor who makes recommendations and stands by to alert the user's attention whenever he is about to commit an error in the interpretation of the data, to ignore important alternatives or to request information - that explicitly or implicitly - is already available. Of course, the user may override any suggestion made by the system, and if he insists on adhering to his "erroneous" decision, the system is loyal enough to stay with him and continue to provide useful analyses and recommendations. At the user's request, the system is also capable of explaining the reasoning process behind its recommendations for either situation assessment or requests for information.

To be able to operate in this manner the system must be so intelligent as to perform by itself the situation assessment task at an expert level. Ideally, we would like it to be able to replace the human decision maker and take over if necessary ("automatic pilot"). Practically, however, we do not foresee a situation where machine alone (or man alone?) can perform better than a human-machine team.

The design and development of an intelligent computer-aided military situation assessment system consists of three major interrelated tasks. The first task is concerned with the elicitation and computer representation of the necessary military knowledge for battlefield reading. This includes, for instance, the characteristics of the various situations in a battlefield and the information sources which are available to the commander and their cost and reliability. The problem focuses on capturing the essence of the situation assessment process and its elements in a fashion which can later be utilized by inference algorithms for situation assessment.

The second task is concerned with modeling the reasoning and inference processes which take part during situation assessment. These include recognizing patterns of military indicators, evaluating and choosing information sources and composing a picture of the battlefield.

The third task is concerned with the human-machine interface, i.e. the design and development of the language and the other means by which a commander dialogues with the system and makes optimal use of its capabilities.

5. REQUIREMENTS ANALYSIS FOR KNOWLEDGE REPRESENTATION AND ELICITATION

Situation assessment is a complex process involving many elements and interactions among the wide variety of battlefield components. A large number of data items are included and used in this process, many of them implicitly. It is not a simple task for a commander to verbalize and spell out the reasoning process which guided him in the analysis of a certain situation. On the other hand, a basic requirement for any intelligent computer system for military command is a systematic and structural representation of military knowledge. The transfer of knowledge from expert human beings to a computer system requires, therefore, two elements. The first is the development of an information structure to accommodate the experts' knowledge. The second is an elicitation technique by which the necessary military knowledge is extracted from expert commanders, manuals, and existing data bases. Of course, the information structure must be designed with the elicitation requirements in mind so that an optimal military knowledge base will emerge.

The elicitation of military knowledge presents unique problems which stem from the fact that recent years have seen very few real large-scale battles. As a result, statistical battle data is not available, and the number of officers with actual battle experience is decreasing. This implies that a military knowledge base can rely only to a limited extent on previous experience. Rather, it will have to rely extensively on subjective and judgmental understanding of the overall doctrine of the opponent. The requirements for a technique for knowledge representation and elicitation include the following:

Compatibility with Human's Cognitive Processes. The most fundamental requirement for any knowledge elicitation technique is the compatibility with the knowledge that a human expert can provide adequately. An example would be a request from an expert to estimate the probability of a given class when a group of features is present (a request which is typical in many rule-based

systems, e.g. [7]). Such a request should be the last resort since it is well known that human is not so good in aggregating adequately the significance of a group of features [18], [5]. Further discussion regarding the ability of human to provide subjective information may be found in [19], [20].

Group Elicitation. To avoid any personal bias, mistakes, or lack of knowledge of a given individual, each component of the knowledge base must be produced by a team of experts. Group elicitation techniques, and their characteristics and requirements, including aggregation methods, have been extensively discussed in the literature, e.g., Huber [21], Linstone and Taroff [22], Keeney and Raiffa [23] and Dalkey [24].

Modularity and Efficient Integration. Because of the versatile aspects of the system knowledge base, its establishment would require several teams of experts, each of which excels in one aspect of the battlefield. It cannot be done by a single team. This implies a modular elicitation approach by which each team is assigned a module within the framework of its expertise. The approach must be uniform as much as possible so that the various teams can easily communicate with each other. It should also provide efficient tools for integration of the modules into one unit, and for conflict identification and resolution.

Elicitation from Existing Sources. A great deal of the required knowledge base may exist explicitly or implicitly in textbooks, field manuals or computerized data bases. A few examples include field manuals for geographical analysis, field manuals for weather analysis, manuals for the enemy forces doctrine, and the TOS computerized data base currently under development. The elicitation techniques should make provisions for utilizing as much as possible this kind of literature and data bases, thus saving duplicate efforts and accelerating the establishment of a high quality knowledge base. A typical example would be computer programs which automatically derives knowledge from computerized data bases. For research in this direction see also [36].

Minimal Burden on Experts. The cooperation of military experts is a key issue for successful establishment of a realistic and comprehensive knowledge base. The elicitation process is, by its nature, a lengthy process that requires significant intellectual efforts. Therefore, the more that is done in the direction of facilitating the process, the higher the chances are to gain cooperation. For instance, a brute force approach by which a human expert is required to list rules and possibly chain them so that they cover all the feasible situations would try the patience of even the most cooperative expert. This difficulty has been recognized by developers of rule-based systems [25].

Ease of Update. It is very likely that a high quality knowledge base will not emerge after the first round of sessions with experts. In order to encompass the entire complexity of the situation assessment process, the knowledge base will have to pass many "tune-up" iterations in which elements of the knowledge base will be modified or deleted, and others will be added. The knowledge base may also require modifications due to changes in, or better understanding of, the opponent's doctrine. The information structure must therefore provide for efficient and effective updating of the system knowledge base. For instance, a change in component A should be automatically checked for all of its possible effects on other components. Implicit changes that are obvious and do not require expert intervention should be automatically inserted. Others should be brought up for the experts' attention.

Computational Efficiency. The knowledge base constitutes the focal point of the system and is frequently consulted. In fact, all the system activities center around the knowledge base. Therefore, efficient representation and storage of the knowledge base is of great importance. This is important not only because of economic considerations, but also because of human factors of man-machine communication. If every reference to the knowledge base requires a significant amount of time, the attractiveness of the system to the user would drop sharply.

6. AN APPROACH FOR KNOWLEDGE REPRESENTATION AND ELICITATION

6.1 Features and Classes

Our proposed approach to knowledge representation is based on two main concepts: features and classes. The term "feature" represents any piece of information related to the battlefield situation, for instance, volume of communication activity, number of tanks in a given force, road conditions, activity in rear area, etc. The term "class" refers to a combination of features that constitutes a well-defined situation in any aspect of the battlefield. For instance, classes for the possible deployment of a unit in the attack are: C1: COLUMNS, C2: LINE, C3: WEDGE, etc. The features that characterize these classes include #1: "dispersion of tanks", #2: "direction of attack", #3: "distances between tanks:", etc. The concept structures are basically hierarchical (see Figure 1); that is, classes at level i are features for the next higher levels $i + 1$, $i + 2$, etc.

At the very bottom level (level 1), the features are specific raw data items, i.e., events or activities observed in the battlefield. Patterns of these features create the classes of level 2, which describe indications regarding the enemy activities or intentions. The indications of levels 1 and 2 become the features for level 3 classes, which are either higher level indications or, in fact, constitute final recognition of the enemy course of action. For instance, in Figure 1, which illustrates this structure, the level 1 feature, such as "size of trees > 6 ft", serve as indications for level 2 classes that describe concealment. The possible CONCEALMENT type of level 2 becomes the features of level 3 classes regarding the TERRAIN. Generally, the features that

characterize the classes in level $i+1$ are not restricted to come from the immediately lower level i . They may come from any lower level below. For instance, "distance between tanks > 60 yds" is a feature for DEPLOYMENT and TACTICS classes at level 4. Also, features at any given level may serve as indications for classes in different perspectives of the battlefield. For instance, TERRAIN features are indicators for TYPE classes and for TARGET classes.

A cost is associated with each feature at level 1 that reflects financial, temporal, and logistic efforts required to obtain the information about that feature. For instance, information that is provided by an in place observing officer would be significantly cheaper than information obtained by a reconnaissance aircraft. This cost is used by the algorithms which generates cost-effective information acquisition proposals. (The airborne observer will probably provide more accurate information than the officer, but costs more).

Logical inter-relationships between features are also recognized and utilized in the situation analysis model. These refer to inter-relationships in which the value of a given feature dictates the value or relevancy of another feature. For instance, if the feature "tank force moving to forward position" is negative, then this automatically implies that features such as "type of tanks", "number of tanks", "configuration of tanks", etc., are irrelevant. Or for instance, if the feature "increased activity in rear areas" is negative then this implies that all the features regarding increased activities, e.g., "intensified traffic of fuel tankers" are negative.

The features may also be categorized into groups that suggest relevancy to a given battlefield. For instance, features relevant to Navy operations are irrelevant to a desert battle far away from any large body of water.

Very few classes may be easily recognized by one or two unequivocal features. Typically it is a combination of a number of features which leads to class recog-

dition. Each of the features by itself may not be so indicative, however, as a pattern they solve the puzzle by providing the evidence which characterizes a class and differentiates it from all other classes.

6.2 Global and Local Knowledge

Two types of knowledge may be distinguished to describe feature/class relationships:

1. Global Knowledge - In a complex real life situation assessment tasks a fundamental issue involves the identification of the relevancy of a given feature to a given class. For this purpose we propose (Section 6.3 below) to represent the knowledge base as a network of relevancy pointers which technically may be defined as a semantic network [26], [27]. Attached to the links of this network are weights by which the strength of relevancy is evaluated and global directions (hypotheses) for situation analysis are determined. The aggregated score of several links pointing at a certain direction may be considered as a rough estimate for its validity, and hence may be used for setting goals (i.e. focusing attention) and prioritizing additional information requests.

2. Local Knowledge - Within each class more specific local knowledge is required in order to assess its validity given a specific set of findings. It seems that in most practical problems multiple representation techniques would be required to characterize adequately the classes by means of their relevant features. Here are some examples:

(a) Patterns of Individual Non-Related Features

The representation of a class pattern in this form is useful under two conditions:

1. Many individual features fairly unrelated are significant for recognizing the class, however, not all of them are required for this recognition. (That is, many legitimate manifestations of this class exist.) For instance, if n binary features are significant for recognizing an ATTACK intention, then a

substantial number of the 2^n combinations may be sufficient to indicate an attack with high confidence.

2. A general rule, e.g. based on Bayes formula, or linear aggregation, may be devised to aggregate the significance of any combination of features.

(b) Patterns with Temporal Relationships

For some classes the sequence in time in which the features are observed are crucial for their interpretation. For instance, the movement of certain forces towards the flanks (X_{15}) may indicate an offensive as well as a defensive intention depending on the order in which previous features were observed. One alternative to represent this kind of knowledge would be by explicit rules, such as: If X_6 and X_9 are followed by X_{15} then enemy attack is expected with probability 0.6. This would require a TIME dimension in addition to AND/OR dimensions typically used in rule-based systems.

(c) Patterns with Structural Relationships

Classes such as those related to deployment of forces are characterized by pictorial and structural information. For this purpose, syntactic and structural representation of class patterns [28], [29] is more appropriate.

In the following section we present a unified approach for global knowledge elicitation and representation. Local knowledge seems to be more specific to the problem domain and will not be addressed here.

6.3 Global Knowledge as a Network of Relevancy Pointers

Our approach to elicit and build the network of relevancy pointers required for representing the global knowledge is class-oriented. Namely, the original* formulation of the knowledge base is as a set of patterns each of which represents the characterization of a given class in the hierarchical

*For purposes of efficient memory management and computation we may restructure the knowledge base for real time use.

structure by means of its relevant features and their weights of significance for recognizing the class. A possible Bayesian approach [10], [34] is to indicate the significance of each feature X_j in a given class pattern C_i by means of two conditional probabilities P_{ij} and \bar{P}_{ij} where:

P_{ij} = the probability that feature X_j is positive given that class C_i is the true class.

\bar{P}_{ij} = the probability that feature X_j is positive given that class C_i is not the true class. Namely, the probability that the presence of X_j is attributed to class(es) other than C_i , e.g. deception.

The P_{ij} value represents the sensitivity (true positive rate) of feature X_j as an item of evidence for class i , while the \bar{P}_{ij} value represents one minus the specificity (false positive rate) of X_j as an item of evidence for class i . In other words, the ratio P_{ij}/\bar{P}_{ij} indicates the odds in favor of class i when a positive result is obtained for feature X_j . Similarly, $1-P_{ij}/1-\bar{P}_{ij}$ indicates the odds in favor of class i when a negative result is obtained for feature X_j .

When the classes for a given aspect of the battlefield are mutually exclusive and exhaustive, only P_{ij} needs to be estimated for every feature X_j that is relevant to class i . The value for \bar{P}_{ij} may be calculated as a weighted average of the P_{ij} 's over the rest of the classes in this aspect of the battlefield. It should be emphasized, however, that our proposed model does not necessarily lead to recognition algorithms which assume that the classes are mutually exclusive and exhaustive.

Features that are not included in the pattern of class i are considered irrelevant for that class. The irrelevance of feature X_j to class C_i indicates that information regarding X_j does not affect the assessment whether class i

is a true class or not. For instance, the feature "sudden increase in communication and electronic activities" may provide very little or no information for diagnosing the strength of a given force. Rigorous definition of irrelevant features and probability computations with them are discussed in [30]. Basically, the paper shows that for any given class-pattern, only the conditional probabilities for the relevant features need to be estimated. This is, in contrast to classical Bayesian models in which the class/feature conditional probabilities need to be estimated for every class/feature combination. The result is a major reduction in the number of probabilities that have to be estimated by experts.

Table 1 shows a typical pattern for the ATTACK alternative when analyzing the overall enemy intentions. The indicator list was taken from a field manual while the probabilities were taken from Table 1 in [31]. The \bar{P} values were calculated to be the average over the courses of actions DEFEND, DELAY, and WITHDRAW. Notice that, in principle, some of the features in this pattern may be classes, which by themselves need to be characterized by a class pattern. For instance, to determine that "extensive artillery preparation" is in effect we need to observe several lower level indicators.

Accordingly, the elicitation process is primarily directed from the class domain to the feature domain. That is, first the class is specified and then, for this particular class, the experts supply the characterizing pattern. This is in contrast to several rule-based models, in which the elicitation process is directed from the feature domain to the class domain, i.e., experts are required to provide rules of the form: if feature 1 and feature 2 and ... feature k are present, then class i exists with certainty P.

The "class to feature" direction is the dual direction to that of actual situation assessment, in which the decision-maker first observes features and then tries to infer their meaning. For elicitation purposes, however, we found that the "class to feature" direction is of greater advantage than the "feature

TABLE 1. A PATTERN FOR ATTACK INTENTION

Class: Attack Intention

<u>Features</u>	<u>P</u>	<u>\bar{P}</u>
Massing of mechanized elements	0.8	0.3
Extensive artillery preparation	0.8	0.4
Artillery position concentrated	0.8	0.2
Concentration of mass toward either or both flanks	0.7	0.3
Location of enemy troops in forward assembly area	0.8	0.3
Location of supply and evacuation installation well forward	0.7	0.3
Increased air reconnaissance	0.8	0.4
Movement of additional troops toward the front	0.8	0.4

TABLE 2

INTERVAL ESTIMATES FOR
QUANTIFYING CONDITIONAL PROBABILITIES

A - Always	$P = .1$
VP - Very Probable	$0.90 \leq P < 1$
P - Probable	$0.75 \leq P < 0.9$
F - Frequent	$0.50 \leq P < 0.75$
S - Sometimes	$0.25 \leq P < 0.5$
R - Rare	$0.10 \leq P < 0.25$
VR - Very Rare	$0 < P < 0.10$
N - Never	$P = 0$

to class" direction for the following reasons which will be exemplified by medical diagnosis.

When a disorder is specified, the frame of reference is well circumscribed. Given a disorder, e.g., acute myocardial infarction, the clinician is required to estimate the frequency with which he observes a specific feature, e.g., diaphoresis, in patients with this disorder. On the other hand, when a feature is specified and the clinician is requested to estimate the probability of a certain disorder given this feature, his first reply is typically, "What else do we know about the patient?" Indeed, physicians do not usually think in terms of the diagnostic values of individual features, since diagnosis is most often based on a pattern of features and not on a single feature. When features are uniquely specific and/or highly sensitive, there may be an exception. In this case, we do encourage the clinicians to provide probabilities of disorders given features. These probabilities can be easily translated into P and/or \bar{P} values [10]. Experiments designed to test some aspects of this issue are reported in [37].

The elicitation process of the knowledge base proceeds in three main stages that may repeat themselves until a convergence to high quality patterns is reached.

Stage 1: Class Characteristics. For a given battlefield aspect, say "TYPE OF ATTACK", all the possible classes are first identified, for example C_1 : Frontal Attack, C_2 : Close Envelopment, and C_3 : Deep Envelopment. With the aid of recent literature and using expert judgement, we list for each class those features that are significant for its recognition. Next, estimates for P_{ij} and \bar{P}_{ij} are obtained to quantify the sensitivity and specificity of a feature j for recognizing class i . These estimators do not have to be point estimators. For practical purposes, interval estimators, (e.g., P_{ij} between 0.10 to 0.20) are sufficient, since correct decisions, which are based on the aggregate evidence conveyed by several observed features are not very sensitive

to non-drastic changes in the P_{ij} and \bar{P}_{ij} values [32]. Table 2 shows an example of a scale that may be used for this purpose. The terminology in this table may also be used when presenting conclusions to the user.

A set of utility computer programs may be devised to assist the experts in developing these patterns. These programs guide the expert systematically through all the steps in class characterization, obtain his answers, and directly generate the internal representation of classes, features and probability estimates. The role of these highly interactive programs is to systematize pattern formulation, to facilitate storing, modifying and retrieving of data; and to smoothe the communication channels between experts. The effort required for the development of these programs is marginal compared to the savings in experts' and analysts' time.

Stage 2: Class Differentiation. Having established the initial patterns for all the classes in a given battlefield aspect, the differentiability of each pair of classes is examined. This may be done by a computer program which computes a discrimination measure [33] between each pair of classes and displays those pairs which cannot be distinguished by their current characterization. The experts are then requested to list the features that differentiate between each pair of classes. This may imply adding new features to each pattern or modifying the P_{ij} or \bar{P}_{ij} values. At this stage, we ensure that any pair of different classes can be differentiated by means of observable features.

The number of pairs for each aspect of the battlefield is $M(M-1)/2$, where M is the number of classes in this aspect. Since M is usually not larger than seven and is almost always less than twelve, the number of all possible pairs is manageable.

Stage 3: Feature Characterization. In Stages 1 and 2 the classes serve as the main frame of reference. Stage 3 concerns the knowledge base from the feature perspective. Using the class patterns, the feature patterns are created

by a computer program. Namely, for each feature, all those classes for which it is significant are listed together with the corresponding P_{ij} and \bar{P}_{ij} values. For each feature, the list is reviewed to verify that all the relevant classes to this feature are included. In the event that a relevant class is missing, the pattern of this class is updated to include this feature as well. The P_{ij} and \bar{P}_{ij} values for different classes are also reviewed, and this may suggest modifications to obtain a more appropriate proportion for the distribution of this feature over its relevant classes. This stage of expansion and refinement would improve if a group of experts is consulted.

The end result of Stage 3 is, in fact, the required network of relevancy pointers which serves the inference algorithm for global analysis.

7. SITUATION ASSESSMENT INFERENCE ALGORITHMS

The cyclic situation assessment process (Figure 2) is driven by a set of algorithms as follows: Once a feature X_j , is observed, all the classes for which it is relevant are identified. These classes may be in different classification schemes, where each classification scheme represents one aspect of the battlefield. For a given relevant class i the evidence conveyed by X_j is integrated with the existing evidence using a global analysis algorithm e.g. [34]. This updates the validity score (e.g. probability) of class i and of the other classes for which class i serves as a feature.

Next, the method of local representation of class i is identified by which the inference algorithm is selected to interpret more rigorously the significance of X_j , e.g. rule chaining or syntactic parser. If the significance of X_j cannot be analyzed separately, additional features may be requested, possibly with an explanation as to why they are needed. Otherwise, the present probability of class i is updated to accommodate the evidence conveyed by X_j . The posterior probabilities of all selected classes are compared with two criteria: VERIFIED and ELIMINATE, e.g. [35]. For those satisfying VERIFIED, we decide tentatively that class i is the true class for the battlefield perspective it belongs to. For those satisfying ELIMINATE, we decide tentatively that class i is not the true class. Once a tentative decision is made for class i , it is not taken into consideration for future goal-setting purposes (see below). However, we continue to update the probability of C_i as more features are obtained that are also relevant to C_i . Such an update will either strengthen the tentative decision or raise suspicions regarding its validity. If such an update elevates the probability of an eliminated class so that the ELIMINATE criteria are not satisfied any longer, or decreases the probability of a verified class so that the VERIFIED criteria are not valid any more, then the corresponding tentative decisions are cancelled. If temporal considerations require that certain actions be taken immediately, then these actions should be taken under the assumption that the tentative decisions are final decisions.

The relevant classes that are neither verified nor eliminated are designated as active classes which require additional evidence for classification. They are subjected to the algorithms for hypothesis generation and goal setting. These algorithms try to compose a complete picture of the battlefield using the tentative decisions on verified and eliminated classes, and the status of the active classes. As a result, one or more hypotheses are generated with regard to the overall picture which are derived from hypotheses with regard to individual classes (or vice versa). Goals are then set to explore selected hypotheses which seem to be more promising or more important to conclude the complete interpretation of the situation.

The active and non active classes associated with the current goal(s) point to lower-level features which have not yet been observed and which are relevant to these classes. These features are then evaluated by weighing their potential contribution to recognizing each of these classes against their cost of testing. As a result of this evaluation, the next features to be tested are recommended to the decision maker. When new features arrive this cycle starts again.

8. SYSTEM STRUCTURE AND HUMAN-MACHINE COMMUNICATION

Figure 3 describes the functional top level building blocks of the situation assessment system. The heart of the system is the Situation Assessment Processor, which generates an integrated interpretation of all the available data and events in the form of a unified situation assessment. The output of the Situation Assessment Processor is formatted by the Summary Generator into a document similar to the Intelligence Estimate Report.

The Military Knowledge Base contains the explicit representation of military knowledge. This information is derived from the literature and from experts and is kept in data structures such as those described in Section 6. These data structures are used by the Situation Assessment Processor to direct its activity of situation recognition and information acquisition. The knowledge base may also contain descriptions of the content of external data bases such as a geographical data base, weapon systems data bases, etc.

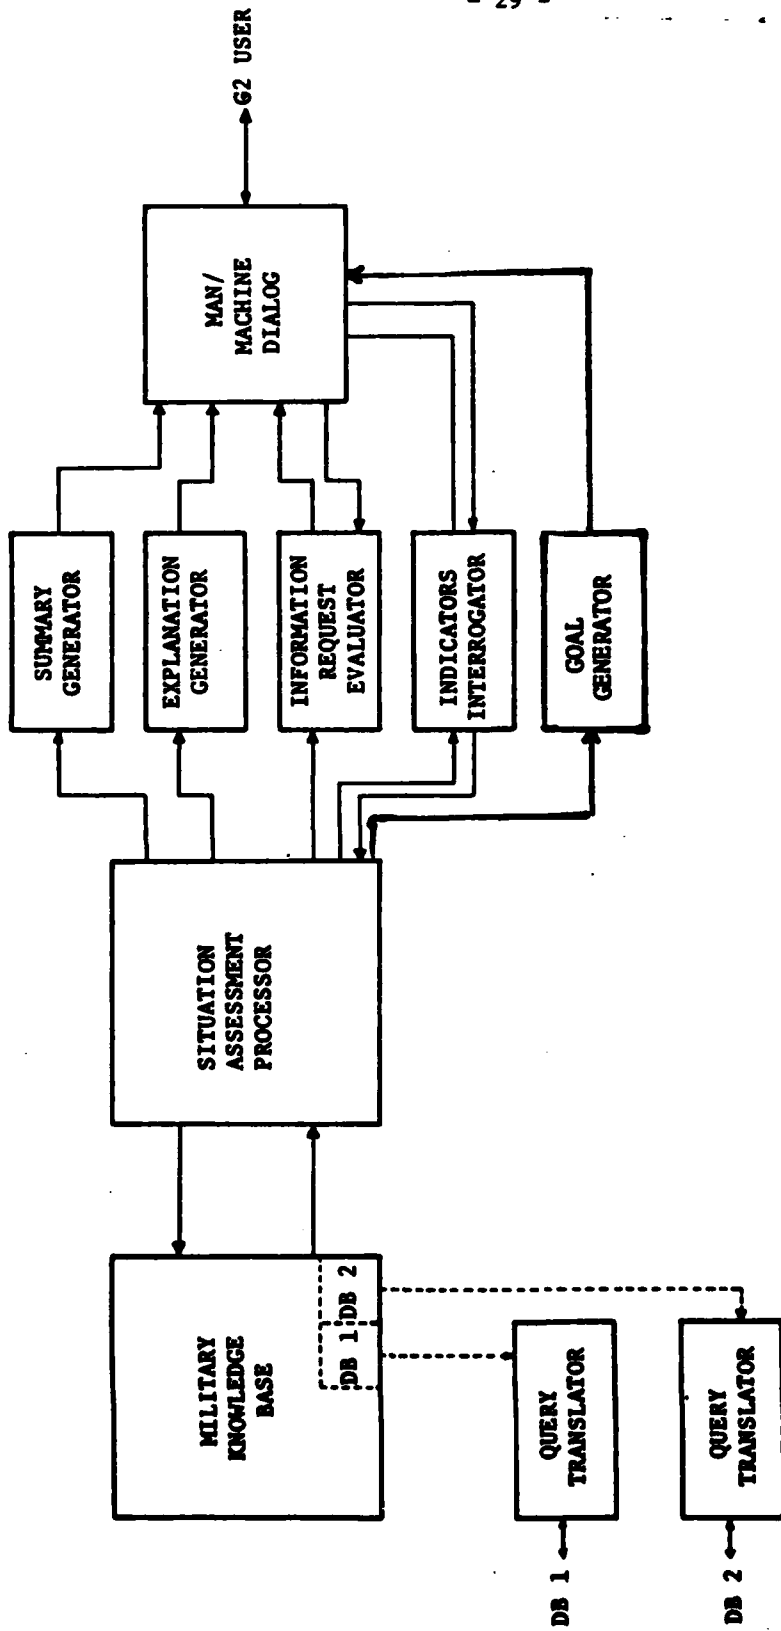


FIGURE 2. SYSTEM BLOCK DIAGRAM

Via these descriptions, the system may address queries directly to these data bases in order to obtain data regarding a particular situation or opponent forces.

The interaction with the system user - the G2 officer - is controlled by the Man-Machine Dialog module. It generates displays, formulates queries, checks input consistency, and handles all user interactions.

Five main types of system-user interactions may be identified, each of which is conducted by the following modules:

The Indicators Interrogator is the module that elicits from the user the facts about the situation and translates them into features as required by the situation assessment processor. The information is elicited through fixed-format queries (e.g., if tanks were observed, the system responds by asking for the number, location, type, activity, etc.). These details are transformed to an appropriate internal representation.

Using predetermined structures (for attack, defense, etc.), the Summary Generator produces a detailed analysis of the situation in light of the chosen interpretation. This analysis is given in the format and structure of the Intelligence Estimate Report commonly provided to the commander by the G2. The summary is centered around enemy intentions and his most probable courses of action.

The Explanation Generator allows the system to produce answers to "How" and "Why" questions issued by the user. For example, "How did the system arrive at given conclusions?" "Why is the system requesting specific information?" "What is the evidence supporting a given interpretation?" This module also permits retrieval of any portion of the system knowledge base for purposes of "on the job" education (e.g., "which are the indications that characterize a certain course of action?").

The Goal Generator conveys to the user the system's tentative opinion of the situation and presents the hypothesis that it would recommend to explore next.

For a given goal(s), the Information Request Generator analyzes the capabilities and costs of the various information collection resources available to the G2, and recommends an information acquisition strategy.

The key characteristic of the man-machine communication in such a system is the flexibility it offers to the user in controlling its operation. The user may specify the perspectives of the battlefield he wishes to consider, or he may let the computer choose them for him. The user may specify the features he would prefer to observe next, or he may let the system select them for him as well. In between the extreme alternatives of full-user control and full-system control, there exists a wide variety of combinations of mixed-initiative user-system cooperation. The principle is to let the user decide on the operation strategy, and to make the system adaptive so that it can adjust to any operation mode.

SUMMARY

Situation assessment tasks share enough common characteristics that warrants addressing them as a general family of problem solving tasks. In this paper, we characterize the generic nature of this family as a multiperspective multimembership hierarcical pattern recognition problem, identify the types of decision problems involved in the situation assessment process, and propose a unified approach for the development of situation assessment decision support systems (DSS). The focus is on knowledge representation and elicitation, although issues related to inference mechanisms, system structure and expert-machine-user interface are also discussed. The presentation is accompanied mostly by examples from military situation assessment. However, comparable examples from medical and business applications are also cited. In fact, many of the ideas presented here have already been implemented in the MEDAS system; a medical DSS for emergency and critical care medicine.

The approach to knowledge representation is based on the observation that two types of knowledge are required; global knowledge and local knowledge. Global knowledge is required to determine directions on which to focus attention, while local knowledge is required for assessing the validity of a specific alternative based on a given set of findings. Global knowledge is represented as a network of relevancy pointers between classes and features. Attached to the links of this network are weights by which the strength of relevancy is evaluated and global directions (hypotheses) for situation analysis are determined. For local knowledge, it seems that in most practical problems multiple representation techniques would be required to characterize adequately the classes by means of their relevant features. Examples include patterns of individual non-related features, patterns with temporal relationships and patterns with structural relationships.

While local knowledge seems to be more specific to the problem domain, the technique described here for global knowledge representation and elicitation is applicable to many situation assessment tasks.

Situation assessment may basically be considered as a puzzle building task. The proposed structure and techniques offer a systematic framework for organizing the knowledge required for solving this puzzle and placing the information items in a manner which facilitates the creation of a meaningful picture of the given situation. Several research projects are now in progress aimed at the various decision problems which are included in the cycle of the situation assessment process.

REFERENCES

1. L. B. Lusted, Introduction to Medical Decision Making. Springfield, IL: Charles C. Thomas, 1968.
2. A. H. Murphy & R. L. Winkler, Forecasters and Probability in Forecasts: Some Current Problems. Bulletin of the American Meteorological Society, 52:239-247, 1971.
3. H. W. Sinaiko, Operational Decision Aids: A Program of Applied Research for Naval Command and Control Systems. Smithsonian Institute, Washington, D. C. Report No. TR-5, July 1977.
4. F. T. DeDombal, D. J. Leaper, J. R. Staniland, et al., Computer-Aided Diagnosis of Acute Abdominal Pain, Brit. Med. J., 2:9-13, 1972.
5. B. H. Beach, Expert Judgment About Uncertainty: Bayesian Decision Making in Realistic Settings, Organizational Behavior and Human Performance, 14:10-59, 1973.
6. T. Y. Young and T. W. Calvert, Classification, Estimation and Pattern Recognition, New York, Elsevier Publishing Company, 1974.
7. E. H. Shortliffe, Computer-Based Medical Consultations: MYCIN. New York: Elsevier, 1976.
8. H. E. Pople, The Formation of Composite Hypotheses in Diagnostic Problem Solving: An Exercise in Synthetic Reasoning, in Proc. 5th Int. Joint Conf. on Artificial Intelligence, Boston, MA., 1977.

9. R. Duda, J. Gaschnig, and P. Hart, Model Design in the PROSPECTOR Consultant System for Mineral Exploration, in Expert Systems in the Microelectronic Age (ed. D. Michie), Edinburgh: Edinburgh University Press, 153-167, 1979.
10. M. Ben-Bassat, R. W. Carlson, V. K. Puri, E. Lipnick, L. D. Portigal and M. H. Weil: Pattern-Based Interactive Diagnosis of Multiple Disorders: The MEDAS System. IEEE Trans. on Pattern Analysis and Machine Intelligence, PAMI-2:148-160, 1980
11. D. Michie, Expert Systems, The Computer Journal, 23:369-376, 1980.
12. C. A. Kulikowsky, Artificial Intelligence Methods and Systems for Medical Consultation, IEEE Trans. Pattern Analysis and Machine Intell. PAMI-2:464-476, 1980.
13. E. H. Shortliffe, B. G. Buchanan and E. A. Feigenbaum, Knowledge Engineering For Medical Decision Making: A Review of Computer-Based Clinical Decision Aids, Proc. IEEE 67:1207-1224, 1979.
14. F. Gomez and B. Chandrasekaran, Knowledge Organization and Distribution for Medical Diagnosis, IEEE Trans. Systems Man and Cybernetics, SMC-11:34-42, 1981.
15. R. Davis et al., Production Rules as a Representation for a Knowledge-Based Consultation Program, Artificial Intelligence 8:15-45, 1976.
16. W. van Melle, Would You Like an Advice on Another Horn, MYCIN Project Internal Working Paper, Stanford University, Stanford, California, 1974.

17. E. de Bono, Lateral Thinking, New York, Harper & Row Publishers, 1970.
18. P. Slovic, B. Fischloff and S. Lichtenstein. Behavioral Decision Theory, Annual Review of Psychology, 28:1-39, 1977.
19. R. M. Hogarth, Cognitive Procedures and the Assessment of Subjective Probability Distributions, J. Amer. Stat. Assoc. 70:271-294, 1975.
20. R. M. Hogarth, Judgement and Choice, New York, John Wiley, 1981.
21. G. P. Huber, Methods for Quantifying Subjective Probabilities and Multi-Attribute Utilities. Decision Sciences, 5:430-458, 1974.
22. H. A. Linstone and M. Turoff (Eds.). The Delphi Method: Techniques and Applications. Reading, MA.: Addison-Wesley, 1975.
23. R. L. Keeney and H. Raiffa. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. New York: John Wiley, 1976.
24. N. C. Dalkey, Group Decision Theory. UCLA School of Engineering, Los Angeles, California, Report No. UCLA-ENG-7749, July 1977.
25. E. A. Feigenbaum, The Art of Artificial Intelligence, pp. 1014-1029, Fifth International Joint Conference on Artificial Intelligence, MIT, Cambridge, Mass., August 1977.
26. W. A. Woods, What's in a Link: Foundations for Semantic Networks., in P. G. Bobrow and A. Collins (Eds.), Representation and Understanding, New York, Academic Press, 1975.

27. R. O. Duda, P. E. Hart, N. J. Nilsson and G. L. Sutherland, Semantic Network Representations in Rule-Based Inference Systems, in D. Waterman and F. Hayes-Roth (Ed.), Pattern Directed Inference Systems, New York, Academic Press, 1978.
28. K. S. Fu, Syntactic Methods in Pattern Recognition, New York, Academic Press, 1974.
29. R. C. Gonzalez and M. G. Tomason, Syntactic Pattern Recognition, Reading, Massachusetts, Addison-Wesley, 1978.
30. M. Ben-Bassat, Irrelevant Features in Pattern Recognition. IEEE Trans. on Computers, C-27:746-749, 1978.
31. E. M. Johnson, The Perception of Tactical Intelligence Indications: A Replication, U. S. Army Research Institute for the Behavioral and Social Sciences, Alexandria, Virginia, Technical Paper 282, September, 1977.
32. M. Ben-Bassat and K. Klove, Sensitivity Analysis in Bayesian Classification Models: Multiplicative Deviations. IEEE Trans. on Pattern Analysis and Machine Intelligence, PAMI-2:251-262, 1980.
33. M. Ben-Bassat, Use of Distance Measures, Information Measures and Error Bounds in Feature Selection. In: P. R. Krishnaiah and L. N. Kanal (Eds.): The Handbook of Statistics, II, North Holland Publishers (in press), 1981.
34. M. Ben-Bassat, Multimembership and Multipurpose Classification: Introduction, Applications and a Bayesian Approach. IEEE Trans. on Systems, Man and Cybernetics, SMC-10:331-336, 1980.

35. M. Ben-Bassat, D. Issers, J. Levy and M. H. Weil, Goal-Oriented Bayesian Diagnosis: Beyond the Posterior Probabilities, IEEE Trans. Pattern Analysis Machine Intell. (Submitted) November, 1981.

36. R. Davis, Knowledge Acquisition in Rule-Based Systems, Knowledge About Representations as a Basis for System Construction and Maintenance, in D. Waterman and F. Hayes-Roth (Ed.), Pattern Directed Inference Systems, New York, Academic Press, 1978.

37. M. Eurns and J. Pearl, On the Value of Synthetic Judgements, Technical Report UCLA-ENG-CSL-8032, School of Engineering, UCLA, Los Angeles, California, 90024, 1980.