TAXONOMY, STRUCTURE AND IMPLEMENTATION OF EVIDENTIAL REASONING MODELS

Moshe Ben-Bassat

February 1987
CSD-870005

TECHNICAL REPORT
R-78
CSD-8700==
September 1986

TAXONOMY, STRUCTURE AND IMPLEMENTATION OF EVIDENTIAL REASONING MODELS* †

Moshe Ben-Bassat

Faculty of Management, Tel Aviv University, Tel Aviv, Israel
&
Cognitive Systems Laboratory, UCLA Computer Science Department, L.A., CA. 90024

ABSTRACT

The fundamental elements of evidential reasoning problems are described, followed by a discussion of the structure of various types of problems. Bayesian inference networks and state space formalism are used as the tools for problem representation. A human-oriented decision-making cycle for solving evidential reasoning problems is described and illustrated for a military situation assessment problem. The implementation of this cycle may serve as the basis for an expert system shell for evidential reasoning: i.e., a situation assessment processor.

Keywords: evidential reasoning, Bayesian inference networks, expert systems, situation assessment

*This work was supported in part by the National Science Foundation, Grant DSR 83-13875;
† presented at the 2nd AAAI Workshop on Uncertainty in AI, Philadelphia, PA., August 1986.

Acknowledgment: The author is grateful to Dr. Judea Pearl for many useful discussions.

TAXONOMY, STRUCTURE AND IMPLEMENTATION OF EVIDENTIAL REASONING MODELS

Moshe Ben-Bassat

I. PROBLEM STATEMENT

Evidential reasoning refers to inference mechanisms by which evidence provided by a set of indicators (e.g., findings, features, attributes, variables) is analyzed in order to gain better understanding of a given hypothesis, concept, situation or phenomenon. Among problem-solving tasks of this nature are medical diagnosis, weather forecasting, corporate assessment, political crisis assessment and battlefield reading. All of these problems share common characteristics in that the problem solver (PS) starts with an initial incomplete understanding of the situation (e.g) patient status) and based on his prior knowledge and expectations, he seeks additional information that may reduce the uncertainty regarding the complete picture of the situation. Following a cyclic process, sources for additional information are identified, evaluated and utilized; the new evidence is integrated into the existing knowledge base; the situation is reassessed; and, if final assessments still cannot be made, further information is requested. The process ends only when:

1. the problem solver (PS) decides that he knows enough to form a defensible interpretation of the situation; or

2. no additional sources of information can contribute significantly (in a cost-effective manner) to remove the uncertainty which still remains; or

3. time constraints force him to abort the information acquisition process and assess the situation as best he can.

Models for evidential reasoning and uncertainty management have attracted significant scientific effort since the beginning of the century. Classical probability theory, Carnap's and Hempel's confirmation theory, Shafer's evidence theory and Zadeh's possibility theory represent a sampling of such endeavor. Many of the classical models have been adopted and improved by artificial intelligence researchers who have applied them in such varied expert systems as MYCIN [Shortliffe 1976], PROSPECTOR [Duda 1979] and MEDAS [Ben-Bassat 1980].

This paper presents a draft taxonomy of evidential reasoning problems and proposes a framework by which evidential reasoning models may be evaluated and compared. As a frame of reference, we propose models based on Bayesian (probabilistic) inference networks.

2. PROBLEM AND KNOWLEDGE REPRESENTATION

2.1 Bayesian Inference Networks

Problem and knowledge representation for evidential reasoning tasks may be based on uncertain hierarchical inference networks. Typically, in such networks, leaf nodes represent *observable* events (indicators), while higher-level nodes represent events (hypotheses) the values for which (true, false or other) may be *inferred* from other (usually lower-level) nodes in the network.

Formally, a node represents a multi-valued proposition in which the values are mutually exclusive and exhaustive. If this is not the case, the node is decomposed into separate nodes, each of which represents a set of mutually exclusive propositions.

At any given time, each node is assigned a set of values corresponding to our degree of belief in the validity of the alternative propositions represented by that node. In Bayesian networks, node values represent the probabilities of the various alternatives corresponding to that node.

A link between nodes H_i and E_j represents evidential relevancy between the two corresponding events. Each link is assigned value(s) that represent the degree of significance for inferring h_i from E_j, or vice-versa. In Bayesian networks, a directed link emanating from H_i and pointing to E_j is assigned a matrix that represents $P(e_j \mid h_i)$ for all the possible values of H_i and E_j. Using this formulation, we avoid committing ourselves to whether the link represents a *causal* relationship (i.e., $P($ *symptom* $\mid$ *disease* $)$ or a *diagnostic* relationship (i.e., $P($ *disease* $\mid$ *symptom* $)$. In our experience, however, eliciting causal probabilities is preferable in most cases. See [Ben-Bassat 1980, p.150] for a discussion in the context of medical diagnosis.

Once an observable node is reported, its evidence is propagated along the network links and revises our belief in the validity of the higher-level hypotheses connected to that node. In Bayesian networks, propagation mechanisms are based on Bayes theorem as the fundamental tool for probability revision, e.g., [Pearl 1986a].

2.2 Node Categorization

The hierarchical network structure suggests a categorization of the nodes into three main types. Typically, leaf nodes represent events that can be perceived directly by the system sensors (the "eyes," "ears" and the keyboard), higher-level nodes represent events that are deduced by the system inference engine ("brain"). Root nodes represent the target hypotheses whose resolution is the ultimate objective of the system. Intermediate nodes may, or may not, be on the list of target hypotheses (see below): in any case, we use them to form defensible argumentations of the resolution of higher level hypotheses.

Several comments, however, are in order:

1. An intermediate- or top-level node may sometimes be *directly* observable, but at a higher cost than *inferring* it from observable lower-level indicators. For instance, opening the abdomen (explorative laparotomy) provides direct observation of events that we intially attempt to deduce from less expensive observations.

2. We may sometimes wish to bypass low-level nodes and report a value directly into an intermediate- or top-level node. This value is not an observation; rather, it is a deduction performed by an autonomous agent who is unable to (or prefers not to) delineate the basis for his deduction by lower-level nodes. An example would be a distributed military intelligence operation in which medium-level officers report upwardly only their summarized assessments.

3. An observable indicator may sometimes be observed with noise. In this case we upwardly report a set of probabilities summarizing our impression of the noisy observation with regard to the possible values of the node. An example would be a patient who does not respond unequivocally to a physician's questions.

4. Although the structure indicates that evidence is propagated bottom-up, top-down and sideways propagation may sometimes prove very useful. In fact, an important feature of Bayesian propagation is that it permits propagation in all directions.

2.2 State Space Representation

Using this framework, we may represent evidential reasoning tasks by a state space formalism. The state of the system at any given stage is characterized by the current values on the network nodes. For the initial state, S_O, we assign all top-level hypotheses their prior probabilities. The initial values for intermediate and leaf nodes may be derived from their parents and the values on the links. Additionally, observable nodes that have not yet been observed are assigned the value UNOBSERVED, designated by "?." These nodes are suggested candidates for direct observation by the information acquisition process.

From the goal-state point of view, the network nodes are divided into target and non-target nodes. Target nodes represent hypotheses requiring resolution by the end of the process, i.e., upon process termination, one needs to commit to set of values -- one for each target node -- that jointly constitute the best explanation for the existing evidence.

The set of target nodes depends on the application. In some cases, the only important decisions made at the final stage concern the root nodes; decisions about other nodes do not lead to any operational consequences. In these cases, the goal state depends only on the values of the root nodes.

In other cases, the states of intermediate nodes impact the action plan (in addition to their roles as mediators for higher-level deductions). For instance, in the medical diagnosis of critical care disorders, a node representing the state of SHOCK is an intermediate node. Yet, to devise a treatment plan, it is very important to know whether the patient *is* or *is not* in SHOCK.

Optimal termination criteria and commitment rules are complex issues, still in their infancies for probabilistic inference networks. See, however, [Ben-Bassat 1980b] and [Pearl 1986b]. An example of a simplified goal state for medical diagnosis is as follows:

S_G The values of the top level hypotheses and a selected group of intermediate hypotheses are above or below certain thresholds.

A more sophisticated state space formulation of diagnostic problems is presented by [Ben-Bassat 1985a].

Evidential reasoning is the process of transferring the network from its initial state, S_O, to a goal state, S_G. The operators for this transformation are queries on the observable nodes. The objective of a control strategy [Ben-Bassat 1985b] is to reach a goal state in a cost-effective manner.

3. TYPES OF EVIDENTIAL REASONING PROBLEMS

Three main factors play a role in determining the difficulty of evidential reasoning problems:

1. - network structure (depth, width, loops ...);
2. - target nodes -- their number and interrelationships among them; and
3. - dependencies among observable nodes.

We shall describe several types of evidential reasoning problems based on the first two factors only. These are graphically illustrated in Figures 1 through 6, where the following notation is used:

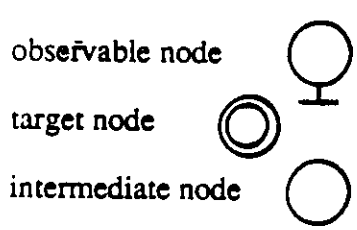

Case (a):

- one set of hypotheses which are mutually exclusive and exhaustive
- observations which are directly linked to the hypotheses

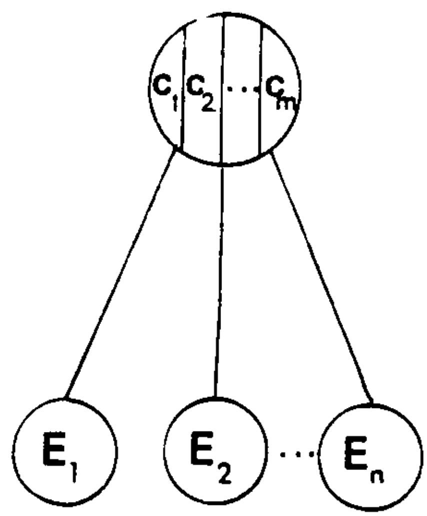

Figure 1 - Classical Bayesian Classification

Case (a) is representative of the well-known classical *Bayesian classification* problem, extensively researched in statistics, decision theory and pattern recognition.

Case (b):

- set of hypotheses which are mutally exclusive and exhaustive
- hierarchical tree-like inferential links

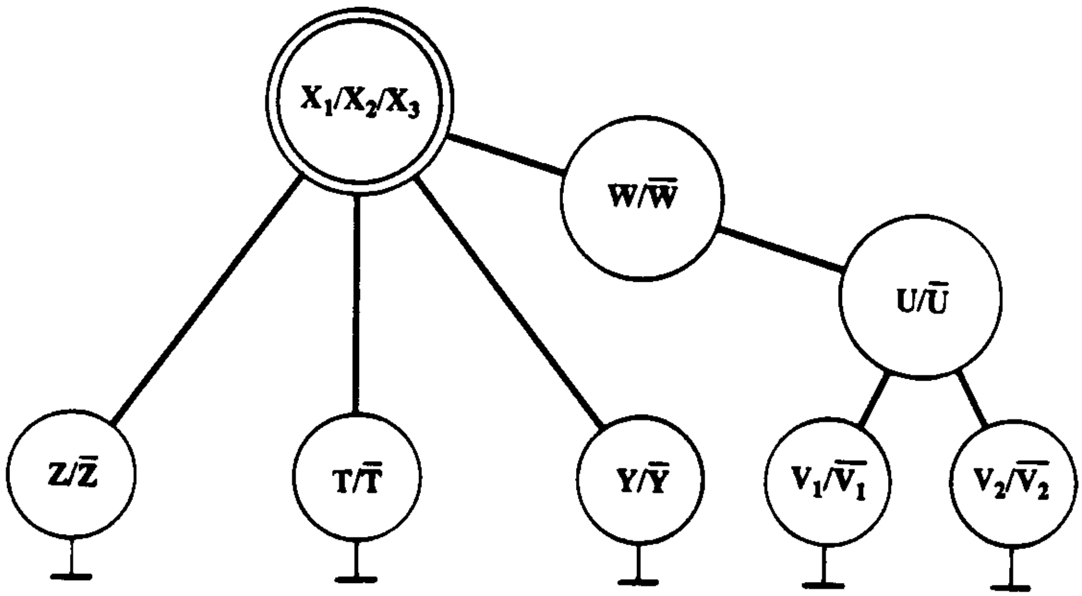

x_1	= HUSBAND (H) KILLED WOMAN (W)
x_2	= HUSBAND DID NOT KILL WOMAN
x_3	= WOMAN COMMITTED SUICIDE
Z	= WEAPON HAS H'S FINGERPRINTS
T	= WEAPON BELONGS TO H
U	= COUPLE HAD MANY FIGHTS
V_I	= NEIGHBOR I SAYS FIGHTS HEARD
W	= H HAS MOTIVE FOR KILLING W
Y	= WOMAN LEARNS TO PREPARE POISON

Figure 2

Researchers in behavioral decision theory refer to Case (b) as *cascaded inference* [Schum 1978]. (The example here was provided by J. Pearl.)

Case (c):

- hierarchically-structured, mutually-exclusive and -exhaustive hypotheses
- observations which are directly linked to groups of hypotheses at different layers of the hierarchy

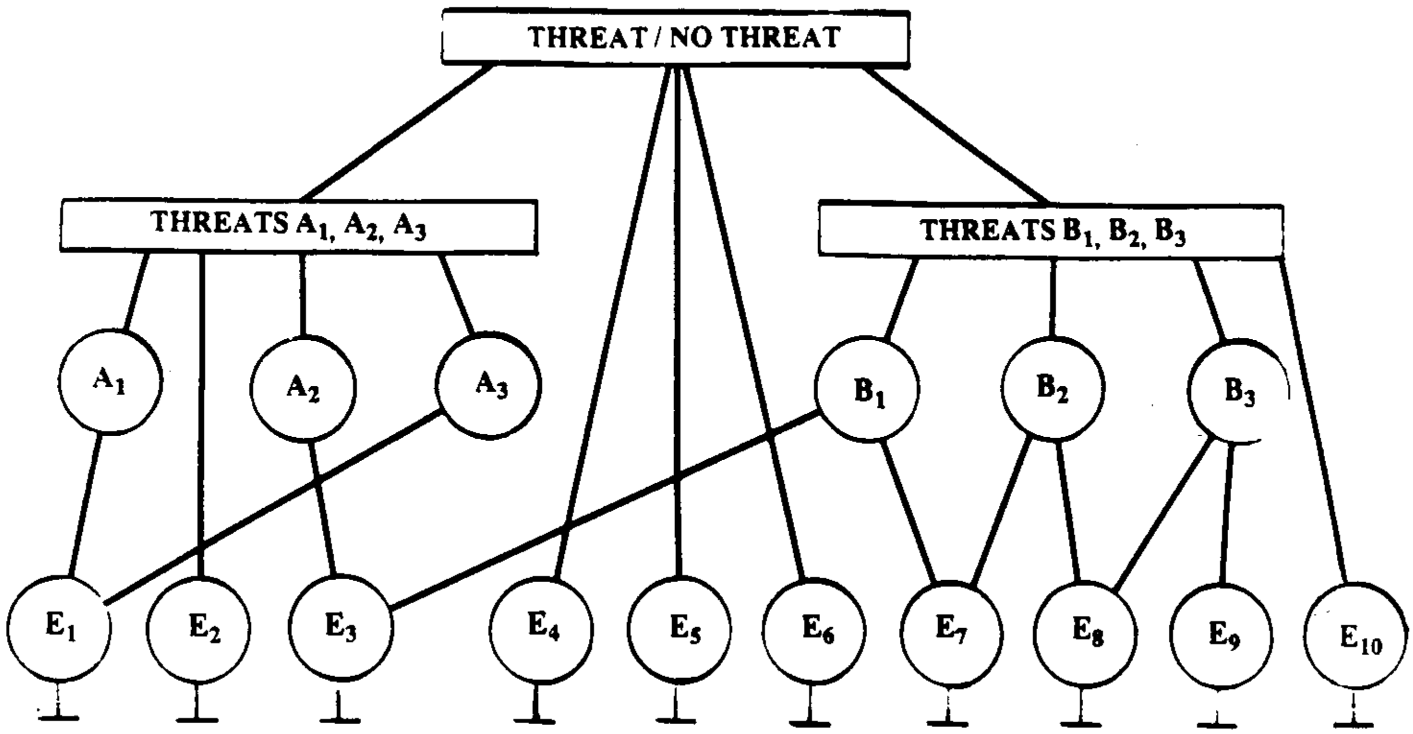

Figure 3

An example of Case (c) appears in *threat assessment of unknown objects*. Some observable nodes may point directly into the root node that represents the hypothesis, whether or not the object is at all threatening. This hypothesis, however, represents several families of hypotheses concerning the various types of threats, each of which may, in turn, be subdivided into more refined classification -- up to the point where each specific type of threat occupies a separate node. For each of the subfamilies, we may have direct links from observable nodes and, perhaps, an intermediate node. Obviously, if E provides evidence for a family H, then it also provides some evidence for all of the subfamilies of H. Gordon and Shortliffe [1985] and Pearl [1986b] have dealt with this problem.

Case (d):

- multiple non-competing sets of hypotheses
- observations which are directly linked to the hypotheses

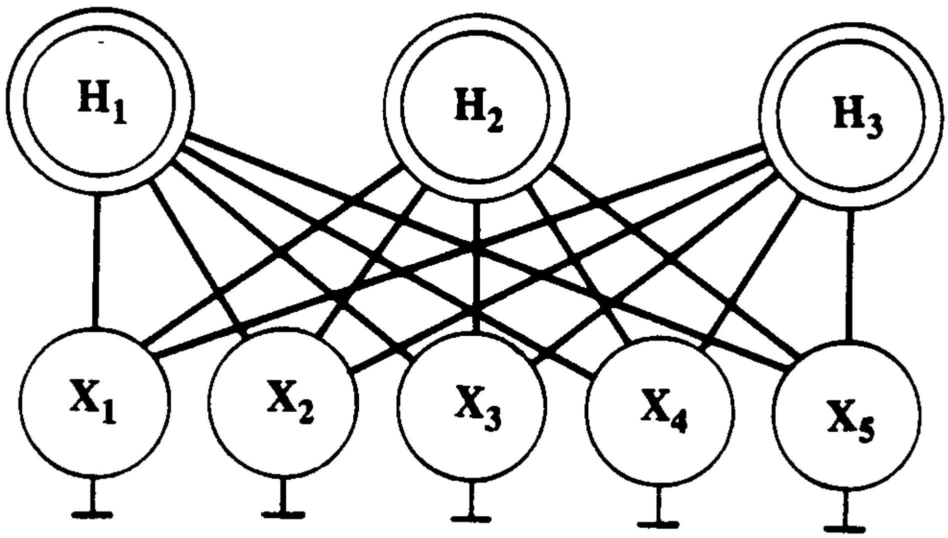

- **A DEFENSIVE FORCE (D) CONTROLS THREE MAJOR POSTS OF A BATTLEFIELD, H_1, H_2, H_3**
- **AN OFFENSIVE FORCE (O) MAY ATTACK NONE, ONE, TWO OR ALL THREE OF THE POSTS**
- **(D) MAY TAKE RECONNAISSANCE ACTIONS IN ORDER TO REDUCE THE UNCERTAINTY. TYPICAL INDICATORS RESULTING FROM SUCH ACTIONS MAY INCLUDE:**

X_1 **- INCREASED ACTIVITY IN THE NORTHERN AREA**
X_2 **- BRIDGING EQUIPMENT MOVED FORWARD**

Figure 4

Case (d) refers to a situation where several hypotheses may co-exist simultaneously. In such a case, every hypothesis H_i gets its own node with two possible values, H_i and $\overline{H}_i$. The problem is also known as *multi-membership classification* in the context of pattern recognition [Ben-Bassat 1980b] and was recently addressed by Pearl [1986c].

Case (e):

- multiple, partially-competing sets of hypotheses
- multi-link inferential chains without loops (singly-connected graph)

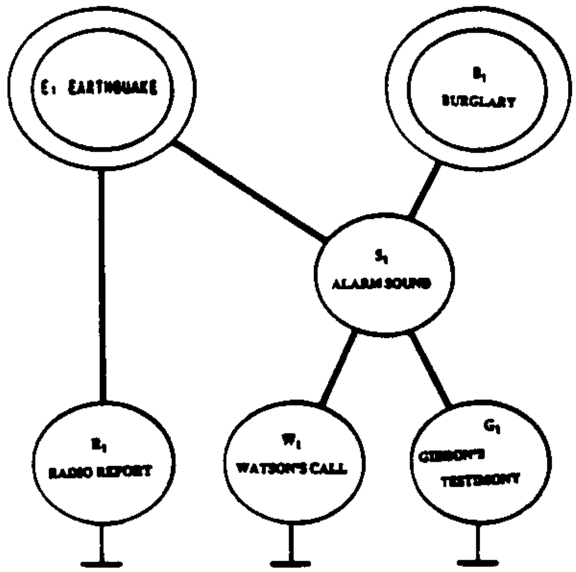

Figure 5

Case (e) represents a situation of multiple causes for a given observation. The example presented in Figure 5 was discussed by Kim and Pearl [1983]. The top-level hypotheses are partially competing in the sense that "earthquake" reduces the likelihood of "burglary" by "explaining away" the alarm sound.

Case (f):

- multi-perpespective hierarchical reasoning (target nodes distributed all over the network)
- no constraints on the links

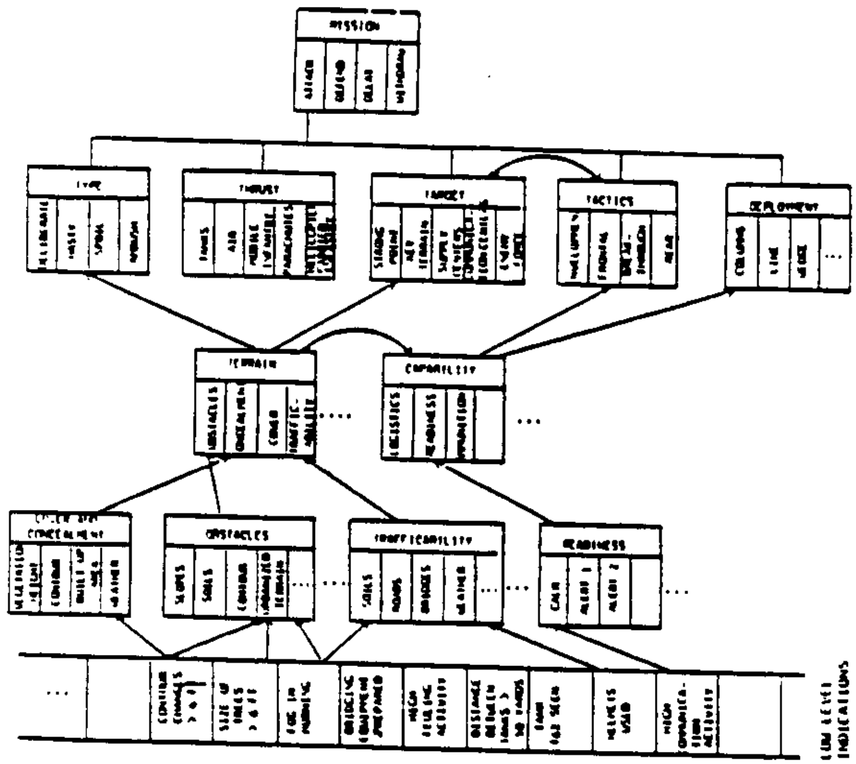

ANALYSIS OF MILITARY SITUATION ASSESSMENT

Figure 6

In many applications (e.g., scene analysis, military situation assessment), in order to generate a rich description of an object or situation, one must view it from multiple perspectives. For instance, in order to analyze a potential military attack, one needs to consider several perspectives [Ben-Bassat & Freedy, 1982]: TYPE, THRUST, TARGET, TACTICS, DEPLOYMENT, etc. Within each of these interrelated perspectives, the situation may be classified by one or more of the alternatives (states or classes) associated with that perspective. For example, an enemy attack can be one of the following TYPES: *deliberate, hasty, spoiling*, or *ambush*. Similarly, there are several alternatives for each of THRUST, TARGET, TACTICS, etc.

Also, within a given perspective, several alternatives may co-exist simultaneously. Within the THRUST perspective, for instance, 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, mobile infantry, parachutes, helicopter-carried infantry, etc.* may, in principle, be simultaneously true.

The recognition process is to some extent "hierarchical" in that low-level indications are used as building blocks for 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 may even influence the choice of a TARGET.

In this case, the target nodes are distributed all over the network because, to devise a battle plan, one needs to know the *details* of both the arena and of the enemy's intentions, not just whether or not he intends to attack.

4. IMPLEMENTATION IN EXPERT SYSTEMS

Expert systems for evidential reasoning problems should support the cycle through which human beings go in solving these problems. Our extensive experience with medical diagnosis, military situation assessment, electronic troubleshooting and other applications suggest the cycle illustrated in Figure 7. Experimental evidence supporting this description may be found in [Eddy and Clanton 1982], [Elstein et al. 1978] and [Zakai et al. 1983].

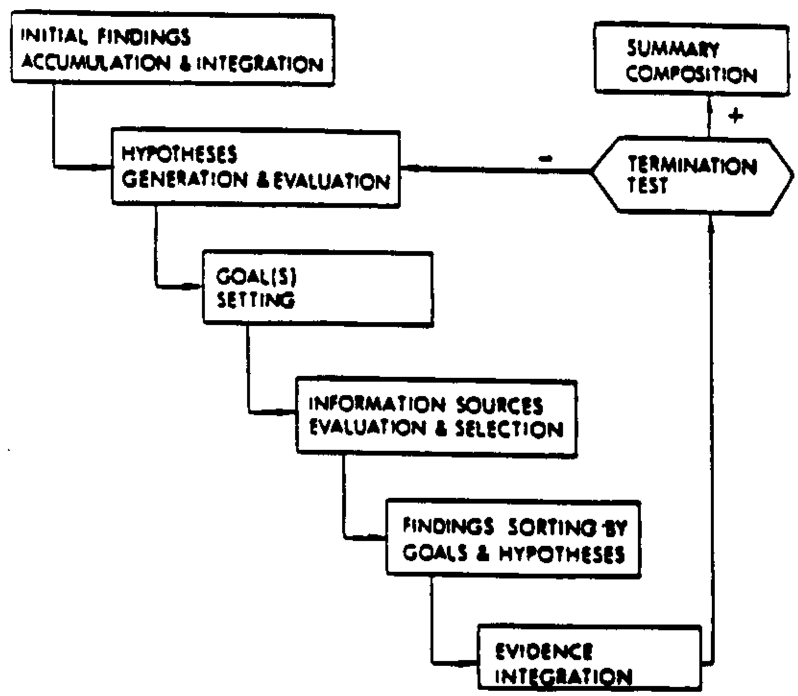

Figure 7 - The Cycle of Evidential Reasoning

Here, for illustrative purposes, a military situation assessment will be used. Each step in the cycle represents one type of decision problem, each of which may require different *skills* on the part of a human problem solver, and different *algorithms* on the part of an expert system.

(0) Accumulating Initial Findings

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 eschelons, from parallel units or from subordinate units. They also include indications or responses to requests for information previously placed 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) Evidence Propagation and Hypothesis Generation

Any recently obtained findings are integrated into the existing body of evidence (which, in the first iteration, is a priori information only) and trigger a moving chain of deductions pointing to several alternative interpretations in several perspectives of the battlefield. The uncertainty regarding the truth of these interpretations is updated and, as a result, some alternatives 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., whether 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 remaining unclear serve as bases for deriving hypotheses to be worked on in subsequent stages. *The generation of a rich set of plausible hypotheses is the hallmark of a good situation assessor.*

(2) Prioritizing Goals Deserving Attention

Occasionally (particularly in early stages), too many hypotheses may be triggered by the existing evidence, and not all of them may be explored simultaneously. In such a case, specific goals need to be set on which we will focus our attention 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 affecting goal determination include the severity and urgency of the candidate alternatives (i.e., enemy attack is expected within 24 hours), their present level of uncertainty and their initial a priori incidence.

(3) Evaluation and Selection of Information Sources

Once a goal is set, information sources potentially offering 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 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) Sorting of Evidence by Goals and Hypotheses

As new evidence comes in, either as a result of the decision maker's request or ''voluntarily,'' it should be sorted with regard to the *entire battlefield structure* (including the triggered hypotheses) and, on the highest priority, with regard to the *current goal(s).* Nonetheless, findings should not be ignored just because they do not contribute directly to the current goal(s) or to the previously activated hypotheses. Such lateral thinking 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 effectively direct the information acquisition path. However, once an indicator is observed, its significance should be analyzed with respect to all of its relevant alternatives.

(5) Evidence Integration

Once all their relevancy links are identified, new findings are *integrated* with the existing findings, *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 may be interrupted or fully terminated under one of the following conditions:

a. A decision is reached with regard to the true situation in each aspect of the battlefield, all (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 gain of information and the impact on the battle plan (or treatment plan if medical diagnosis is the case).

c. New developments (e.g., sudden enemy attack) force the decision maker to abandon information acquisition and, based on existing evidence, assess the situation as best he can.

(7) Integrated 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 (earlier referred to as "commitment decisions"). The end result is the Intelligence Estimate document, which is currently produced manually by the intelligence officer.

5. SUMMARY

We have presented several types of evidential reasoning problems and a detailed description of the Bayesian Inference Networks (BIN) approach for structuring these problems. The references cited in the paper provide a partial picture of the state of the art in Bayesian evidential reasoning. Although much work remains to be done, recent developments and experience in this field suggest that the BIN-based approach is a powerful tool for practical expert systems.

REFERENCES

Ben-Bassat, M., Carlson, R. W., Puri, V. K., Lipnick, E., Portigal, L. D. & Weil, M. H., "Pattern-Based Interactive Diagnosis of Multiple Disorders: the MEDAS System," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, PAMI-2:148-160, 1980a.

Ben-Bassat, M., "Multimembership and Multiperspective Classification: Introduction, Applications and a Bayesian Model," *IEEE Transactions on Systems, Man and Cybernetics*, SCM-10:331-336, 1980b.

Ben-Bassat, M. & Freedy, A., "Knowledge Requirements and Management in Expert Decision Support Systems for (Military) Situation Assessment," *IEEE Transactions on Systems, Man and Cybernetics*, SCM-2:479-490, 1980b.

Ben-Bassat, M., "Expert Systems for Clinical Diagnosis," in Gupta, M. M., Kandel, A., Bandler, W. & Kiszka, Y. B. (Eds.), *Approximate Reasoning in Expert Systems*, North-Holland, Amsterdam, the Netherlands, 671-687, 1985.

Ben-Bassat, M. & Maler, O., "Control Strategies in Uncertain Inference Networks," *Proc. Uncertainty in AI Workshop*, UCLA, Los Angeles, August 1985.

Duda, R., Gaschnig, J. & Hart, P., "Model Design in the PROSPECTOR Consultant System for Mineral Exploration," in Michie, D. (Ed.), *Expert Systems in the Microelectronic Age*, Edinburgh University Press, Edinburgh, 153-167, 1979.

Eddy, D. M. & Clanton, C. H., "The Art of Diagnosis: Solving the Clinicopathological Exercise," *Engl N, Med. J.*, 306:1263-1267, 1982.

Elstein, A. S., Shulman, L. S. & Sprafka, S. A., "Medical Problem Solving: an Analysis of Clinical Reasoning," Harvard University Press, Cambridge, Massachusetts, 1978.

Gordon, J. & Shortliffe, E. H., "A Method of Managing Evidential Reasoning in a Hierarchical Hypothesis Space," *Artificial Intelligence*, 26:323-357, 1985.

Kim, J. & and Pearl, J., "A Computational Model for Combined Causal and Diagnostic Reasoning in Inference Systems," *Proceedings of IJCAI-83*, Karlsruhe, F.R.G., 190-193, August 1983.

Pearl, J., "Fusion, Propagation and Structuring in Belief Networks," *Artificial Intelligence*, Vol. 29, No. 3, September 1986a, pp. 241-288.

Pearl, J., "On Evidential Reasoning in a Hierarchy of Hypotheses," *Artificial Intelligence, 28:9-15, 1986b.*

Pearl, J., ''Distributed Revision of Belief Commitment in Multi-Hypotheses Interpretation,'' Proceedings, 2nd AAAI Workshop on Uncertainty in Artificial Intelligence, Philadelphia, PA., August 1986c, pp. 201-209; also to appear in *Artificial Intelligence*, North-Holland, 1987.

Schum, D. A., "Current Developments on Cascaded Inference Processes," in Wallsten, T. S. (Ed.), *Cognitive Processes in Choice and Decision Behavior*, Erlbaum, Hillsdale, New Jersey, 1981.

Zakai, D., Teeny, D. & Ben-Bassat, M., "Behavioral Strategies in Pattern Classification," *Omega*, 11(4), 385-394, July 1983.

....