Naval Command,
Control and Ocean
Surveillance Center RDT&E Division San Diego, CA
92152-5001

4

Technical Report 1558
December 1992

Development of Computer Support for Naturalistic Decision-Making

D. C. Hair
K. Pickslay

Technical Report 1558
December 1992

Development of Computer Support for Naturalistic Decision-Making

D. C. Hair
K. Pickslay

NAVAL COMMAND, CONTROL AND OCEAN SURVEILLANCE CENTER
RDT&E DIVISION
San Diego, California 92152-5000

J. D. FONTANA, CAPT, USN
Commanding Officer

R. T. SHEARER
Executive Director

ADMINISTRATIVE INFORMATION

This report describes work performed between March 1991 and September 1992. It is part of an ongoing project in the Decision Support and Artificial Intelligence Branch, NCCOSC RDT&E Division Code 444. The basic goal of the project is to explore new approaches for supporting Navy tactical decision makers. The work is sponsored by the Office of Naval Technology.

Released by
R. III. Delgado, Head
Decision Support and AI
Systems Branch

Under authority of
J. D. Grossman, Head
Command Support
Technology Office

RV

EXECUTIVE SUMMARY

OBJECTIVE

Develop software tools that can assist in making decisions in real time environments involving incomplete and uncertain data. Explore ways to incorporate artificial intelligence (AI) techniques into such tools. Focus on support of naturalistic human decision making strategies.

RESULTS

A tool called Situation Assessment By Explanation-based Reasoning (SABER) has been developed to function as one part of a decision support system. SABER makes use of an Explanation-Based Reasoning (EBR) technique to construct and evaluate explanations that account for data which may be incomplete or uncertain. The tool is capable of real time operations, is easily modifiable by domain experts who are not computer experts, and can present a variety of informational displays as desired by users.

RECOMMENDATIONS

1. SABER needs to be tested with domain experts to help refine the user interface and to better explore how it can be of most benefit to users

2. Forthcoming research results from other participants needs to be incorporated in terms of deciding how to best handle the format and timing of displays.

3. Interactions need to be worked out with other tools to be used in the overall decision support system.

CONTENTS

INTRODUCTION

A computer tool called Situation Assessment By Explanation-based Reasoning (SABER) is being developed to assist naval decision makers involved in real-time, tactical, decision-making situations. The SABER work is being done as part of the Tactical Decision Making Under Stress (TADMUS) project, funded by the Office of Naval Technology. SABER makes use of an Explanation-Based Reasoning (EBR) technique to analyze data and interact with users.

Problems arise in tactical situations when data needed for decisions are incomplete or uncertain, and the overall volume of potentially relevant data is high. These characteristics cause problems for human decision makers, and the problems are exacerbated when there are severe time constraints.

The problems are of two basic types: (1) there are difficulties in being able to process the available information quickly enough, and (2) there can be problems with cognitive difficulties. While it is a straightforward proposition that computers can assist in processing data quickly, there has been little work in the area of directly using computers to assist in overcoming cognitive problems.

The TADMUS project is generally concerned with exploring the kinds of biases or other cognitive problems that arise in tactical situations, and with trying to find ways of overcoming such problems. The two basic areas being looked at as sources of improvement are in training and in computer decision-support systems. The SABER work is specifically focused on producing a decision support tool to improve the human decision-making process. It is not intended to replace the decision-making processes of the human decision makers. Improvement is looked for in terms of being able to provide desired information at key times to counteract undesirable cognitive processing patterns.

Other work in the TADMUS project is involved with identifying particular cognitive decision-making problems, and with determining the precise kinds of situations that bring such problems into play. The information generated from that research will be used, in part, to determine how to make best use of tools like SABER.

A key element of the problem area involved in the SABER work is that it is necessary for people to handle reasoning with uncertainty in time-constrained situations. These are situations when there may not be correct a decision, but it is desirable to achieve as good a decision because the consequences of mistakes can be serious. The SABER tool is expected to assist in the decision-making by supplying some of the reasoning at a much faster rate than humans can do. It is also expected that, as a result of using the EBR approach, the tool can structure interactions with users in ways that can overcome some of the problems associated with cognitive biases. In addition, a major emphasis has been put on developing a tool that can have its database and actual results easily modified by decision makers who may not be experts either with computers or mathematics.

The SABER tool models one of the strategies that people use themselves in reasoning with uncertain or incomplete data. The initial impetus behind this work was to implement a version of the EBR technique as a means for reasoning with uncertainty (Pratt, 1987). The EBR approach is justified both because it can function in a variety of well-defined decision-making situations, and because it reflects human decision-making processes in those situations. SABER is expected to interact through a blackboard architecture with other decision-support tools that will comprise a complete decision-support system.

BACKGROUND

There are often suggestions when things go wrong in technological fields that human error is the underlying cause of the problems. That kind of determination is reinforced by research that indicates that people are prone to make a wide variety of incorrect decisions, with suggestions then arising that perhaps computers can be used to supply more analytical approaches that will overcome the weaker heuristic methods used by humans (Kahneman, Slovic & Tversky, 1982).

There has been a more recent trend, however, suggesting that it is quite likely that system design itself often leads to a kind of inevitability of failure at some point, where the root cause of the failure lies in the system rather than in the human users (Perrow, 1984; Norman, 1988). It is also being questioned whether the human decision-making process, in general, is as flawed as some of the research on biases and the like has suggested (Hammond, et al., 1987). A resulting general conclusion is that, humans probably do a good job of decision making which can possibly be made better by designing better systems. The quality of decision making is thought to be particularly high among expert decision makers. Nonetheless, it is believed that some kinds of cognitive biases, or other problems of human cognitions, do arise in a variety of decision-making situations. The SABER work is aimed at trying to support the decision-making process by providing information, together with some preliminary analysis, in a way that can overcome potential cognitive difficulties.

A variety of approaches are possible in the use of computers to support decision makers. One approach is to have the computer use methods that claim to have formal validity, but do not claim to be related to ways in which humans solve problems themselves. Another approach is to not try to use the computer as a reasoning device at all, but instead, to manipulate the computer's abilities to present data through a graphical user interface as a means of influencing the human decisions. A third approach is to model the way that the computer reasons after the manner that humans are thought to reason. Of course, there are various hybrid approaches, but we discuss these as separate approaches to highlight critical features of each.

The more formal approaches tend to rely on extensive calculations designed to arrive at optimal solutions. Recent work has particularly focused on the use of probabilities, using either a Bayesian approach (Charniak, 1991) or the Dempster-Shafer theory (Shafer, 1976). While use of these methods has been quite fruitful, there are two problems that bear directly on the goals of the SABER work. First, these formal approaches tend to require at least exponential time so that the real value of the approaches is in situations where time is not an essential constraint. Second, although proponents of these approaches have made strides in making their resulting tools easy to use, it does not appear that a true layman could successfully set up or modify such systems without expert assistance. The chief advantage of these approaches is that the end result will be, in some sense, optimal when given sufficient time.

The second approach is basically a man-machine interface (MMI). An advantage of this approach is that there is little extra computation time since the pure approach is only looking at ways to manipulate the interface. There may still be time problems, however, if an interface requires a lot of complicated graphics presentations. A drawback is that the full potential for computers to be of assistance is not going to be realized if no use is made of computer reasoning techniques. Of course, a lot of the current work involved with human-computer interfaces takes a hybrid approach by building some kinds of artificial intelligence into the interface.

The modeling approach has the following possible benefits: (1) people may be able to work more easily with systems that use their own natural reasoning strategies, (2) the computer itself is not subject to the kinds of biases that humans are subject to except as those biases may be encoded in programs, and (3) the modeling of human reasoning using a computer can lead to insights about the human reasoning process. Among the techniques that can be classified as modeling are case-based reasoning (Kolodner, 1991), and explanation-based generalization (Mitchell, Keller & Kedar-Cabelli, 1986).

Depending on the process being modeled, tools using the modeling approach can be even slower than those using the formal approaches. The SABER work however has focused on a heuristic strategy that has been implemented in a linear time program.

The relative strength of the SABER approach compared to the formal or MMI approaches can be assessed along three dimensions: computational speed, optimality of results, and ease of use. Because SABER operates in linear time, it is faster than typical tools using the formal approaches. It is likely to be somewhat slower than many of the tools using a strict MMI approach, since those tools typically do not do any analysis of data.

On the dimension of optimality, SABER will not produce results that have the same degree of precision as can be done with the formal approach. The MMI approach is low on the scale of optimality in the sense that it does not make use of the computer itself to compute any optimal results.

In general, it appears that tools based on the modeling approach are easier to engineer for ease of use, including user modifications, than tools based on the formal approach. Thus, tools using the formal approach usually cannot be modified unless the user understands the underlying probability or other theory, whereas, the underlying theories in tools using the modeling approach will tend to be intuitively understood by a wide range of users. Tools using the pure MMI approach usually focus on interface changes that are not modifiable by users at all.

This discussion is not intended to suggest that the modeling approach is better than the other approaches. What is suggested is simply that each of the approaches has its own strengths and weaknesses, and that further study is needed to determine precise circumstances that call for the use of one or another of those approaches.

Similarly, the EBR strategy is not proposed as the only human heuristic strategy that can be modeled in efforts to construct decision support systems. It is recognized that other strategies are used and that support for them is desirable. Other strategies such as recognition-primed decision making are being explored by other researchers in the TADMUS project (Klein, 1989).

In addition to these arguments about the possible strengths of tools based on a modeling approach, it can be noted that there are also recent suggestions in the literature that in some areas it may actually be a necessity for tools to model users' cognitive behaviors (Smith & Lansman, 1992). The goal is to improve a user's performance at an open-ended kind of task, the tool can only accomplish the goal if it can model the user's underlying cognitive behavior. Open-endedness arises in our situations due to the possible presence of ambiguous data.

SPECIFIC PROBLEMS ADDRESSED

The basic problem addressed by this work is to try to improve the decision making processes of people who must make time-constrained decisions based on incomplete and uncertain information. To be useful in such situations, tools must not only try to perform some useful analysis, but must also be able to operate in real time.

Besides being able to reason with uncertainty in real time, decision-support tools should be clearly understood by the users. Problems arise with these tools when users do not see what the tool is intended to do and how it does it.

A related problem to be guarded against is to ensure that this kind of tool is not seen by human users as a decision-making entity in itself, but rather as only an aid to forming their own decisions. Thus, in reasoning with uncertainty we believe it is important that the user understands that the decision-support tool does not necessarily reach correct conclusions, and that the user understands in what ways the tool's analysis can be incorrect.

Another related problem involves the modifiability of the tool. Here again, failure to understand the tool can cause problems with the user's ability to modify the tool in desirable ways. The inability to modify the tool can lead, in turn, to problems in maintaining the usability of the tool. The SABER work has attempted to deal with this kind of problem by enabling users to make changes in virtually every aspect of the underlying knowledge base. Changes are made by having users make choices based on natural language cues rather than on numerical cues.

EXPLANATION-BASED REASONING

There is evidence that in some decision-making situations, people will consider alternative explanations that can account for a given set of data, and will base a decision of the most likely cause of the data by assessing the relative strengths of the explanations (Pennington & Hastie, 1988). The strength of an explanation is based on judgments about the explanation's plausibility. An explanation in this framework is simply a causal model that incorporates all available data into a coherent structure that supports one of the possible decision outcomes.

SABER makes use of related ideas first explored by Hirst (1988). The tool begins by constructing a set of explanations pointing to different conclusions. Explanations are then evaluated according to three primary criteria: simplicity, completeness, and data significance. Taken together, those criteria constitute SABER's measure of plausibility.

Data significance is considered through a weighting scheme. Each type of data is weighted so that the relative importance of different kinds of data can be considered. Two other kinds of weights are also considered. First, when contradictory data are involved, weights assigned to the different assumptions used for explaining contradictions are used to indicate the strength of the user's belief in the assumptions. Second, when certain kinds of data are expected to occur in support of a given conclusion, but have not been observed, negative weights are applied to decrease the degree of belief in that conclusion.

In applying the plausibility criteria, the different explanations are first evaluated by looking for the simplest explanation. Here, preference is given to explanations that do not require elaborate reasoning to explain the presence of data that contradict a given conclusion. For instance, when all available data point to the same conclusion, SABER will normally pick the explanation for that conclusion as the most likely. However, the tool will construct explanations for the other possible conclusions even in that case. Second, SABER looks at how completely a given explanation directly accounts for all available data. In general, if two pieces of data are accounted for by explanation A, and only one piece is accounted for by explanation B, explanation A will be ranked more highly. Lastly, the tool accounts for the differing significance of various pieces of data. For instance, a visual identification of an aircraft may carry more weight than some electronic indications that point to other identifications. Figure 1 illustrates how simplicity and completeness are used.

In addition to those three plausibility criteria, SABER takes into account direct user input in determining the relative strength of the explanations. Thus, users can specify which explanations are most plausible in particular situations, and SABER will use that input in making future determinations.

Before SABER can construct the explanations, a set of possible conclusions and a set of possible kinds of input data must be specified. Each specific kind of data is expected to directly support reaching only one of the possible conclusions. The representation of each data type is itself a set of explanations. The primary explanation associated with a data item indicates how the given kind of data supports one

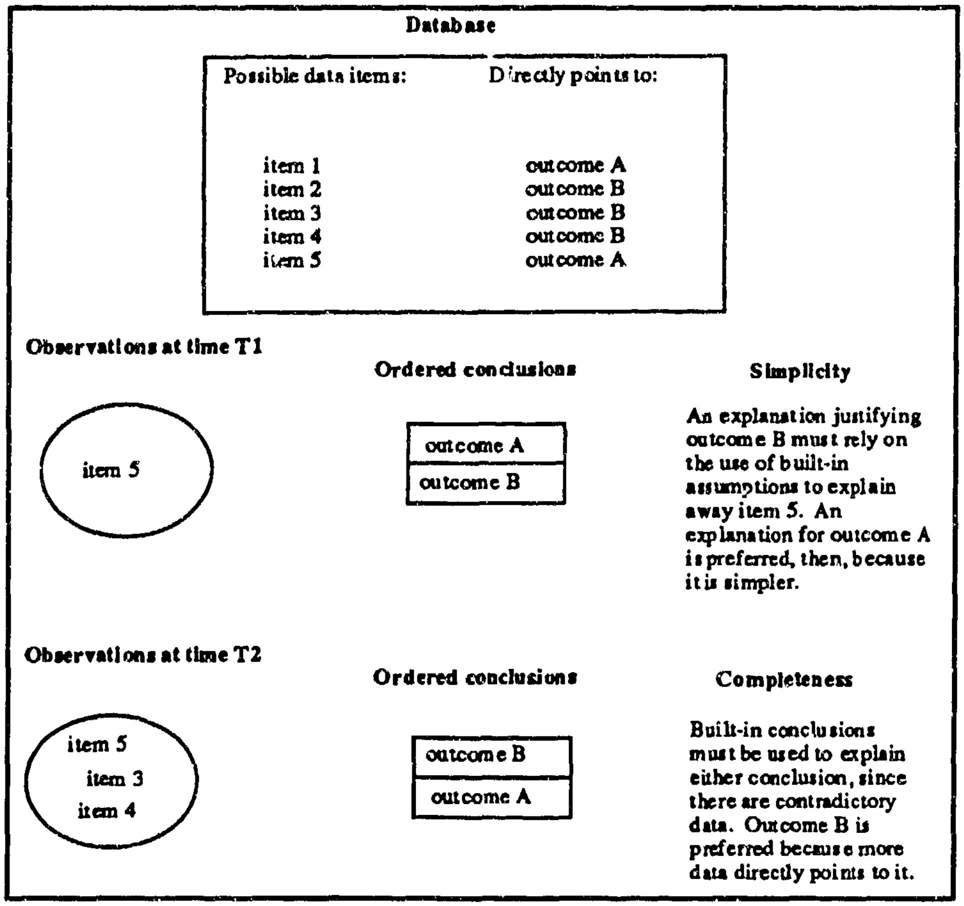

Figure 1. Explanation evaluation criteria.

particular conclusion by default. Alternative explanations are attached to each data item to ndicate the assumptions that could be used to override the default conclusion. See figure 2 for an illustrative representation of a typical data item.

When several kinds of data are present, the individual explanations for each particular piece of data are combined into larger explanatory structures. A larger structure contains separate explanations for each possible conclusion. That structure can then be shown to the user in a variety of ways. Each separate conclusion is explained by presenting (1) the data items that directly support that conclusion, (2) an indication of the strength of the conclusion, and (3) list of the assumptions that can be used to account for the presence of contradictory data. The strength of the most plausible conclusion is indicated by assigning the conclusion one of the following confidence levels: confirmed, probable or possible. Where data are contradictory, the explanation gives assumptions that can be used in determining which one of the possible conclusions should be preferred. (See figure 3.) Assumptions are given their own fuzzy confidence levels, and those levels can be directly modified by the user. SABER uses both the presence and the absence of certain types of data in constructing explanations.

Data type: friendly radar signal

Weight: 48

Directly supports: conclusion that aircraft is friendly.

Reason: friendly radar signals are readily recognized and are expected from friendly aircraft.

Assumptions to support contrary conclusions:

Aircraft could be hostile if:

Hostile aircraft has the capability to mimic friendly radar, or

Radar signal was incorrectly identified due to poor detection conditions.

Aircraft could be neutral if:

Radar signal was incorrectly identified due to poor detection conditions.

Figure 2. Representation of a data type.

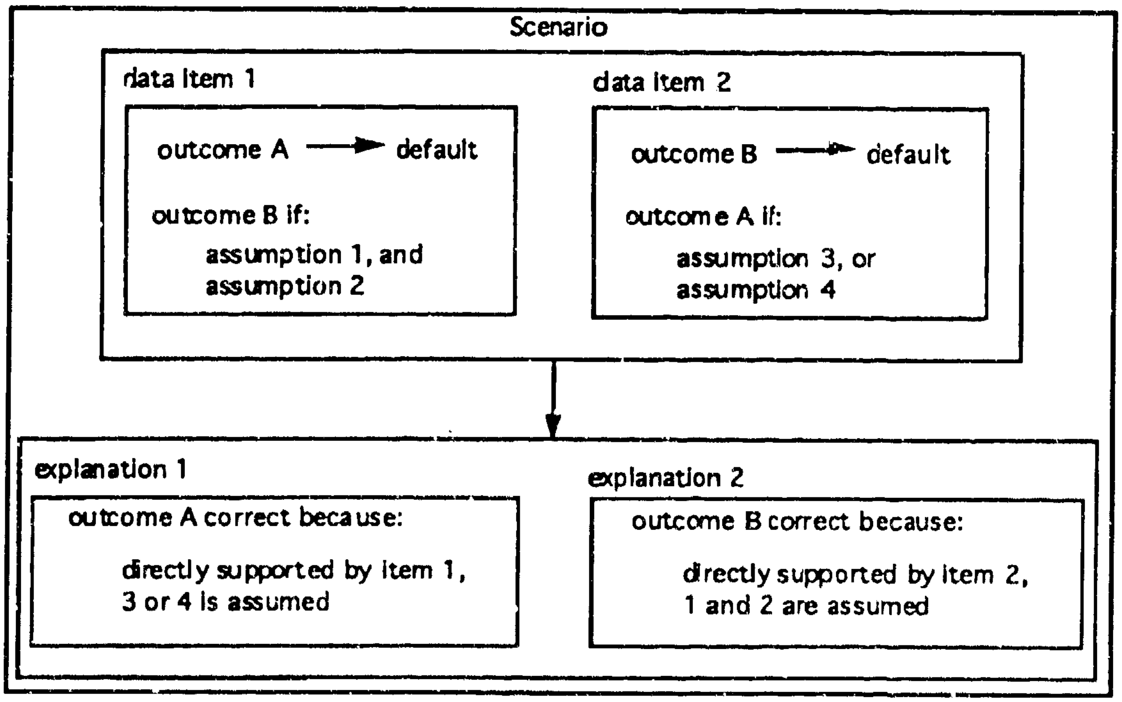

Figure 3. Generation of combined explanations.

SABER has been generalized so that it can construct explanations in any scenario that involves a fixed set of possible kinds of input data and a fixed set of possible conclusions. The primary domain looked at is one in which a user needs to determine whether a given aircraft is friendly, neutral, or hostile based on a set of input indicators, such as, radar or visual identifications. A related possible domain is to determine the type of an aircraft.

USER MODIFICATIONS

In addition to the basic decision-making problem, the SABER work is aimed at producing a tool that will not readily become obsolete. A problem with some knowledge-based tools results from the fact that they are developed with a specific problem in mind and that they [illegible] a specific source of expertise. Obsolescence can result from changes in the nature of the problem or from the expertise becoming outdated. SABER [illegible] designed to be useful in a variety of decision-making situations, and to be easily modified by the end user to add new knowledge or modify existing knowledge.

For the most part, user modifications are intended to be made offline, at times when SABER is not in active use as a decision aid. Thus, it is not believed to be generally appropriate to try to make significant changes in the tool's reasoning or database at times when quick responsiveness is critical.

Users can directly change the evaluations of explanations in three ways. First, users are given the option of specifying which conclusion should be given the highest ranking for any set of possible input items. SABER will then let the user know how the weights given to the data items can be changed to achieve the desired result. In this way, users are allowed to effectively train the tool to reach correct conclusions. Users are not asked to directly specify the weight changes, and are not expected to know the inner workings of the weighting system. The user simply specifies which conclusion to favor and the class of weights to change.

Over a period of time, the overall set of weights is expected to settle so that the tool will always generate preferred conclusions according to an extended test set of situations evaluated by expert users. The idea is somewhat like what is done in neural nets, since here an entire set of internal weights may be changed each time a user indicates that the tool has not reached a desired conclusion, and the weights are expected to reach a near optimal setting. It is expected that the desired setting will be achieved due to anticipated consistency in the expert use of data items to support conclusions. If there is no such consistency, then the results will point to specific problems in this particular decision-making process that will warrant further study.

A second option is for the user to change just the confidence level associated with the preferred explanation of a given set of input data. Here again, weights will be changed to achieve the desired result so that a kind of training takes place. Changes in the confidence level do not affect the choice of the preferred explanation.

Third, users can affect the results by changing the treatment of the assumptions that are used to explain the presence of contradictory data. For instance, SABER might conclude that the best explanation points to a conclusion that an aircraft is hostile even though a signal from the aircraft indicates it is friendly. The explanation of the hostile conclusion might indicate that the friendly signal is assumed to be falsely generated by a hostile aircraft. If the user does not believe a hostile aircraft could be imitating the friendly signal, the user can override that part of the hostile explanation. All of the explanations will then be re-evaluated, and the result may be that the preferred explanation will then suggest a friendly aircraft. This method is the only one of these three intended for use while SABER is in active use, and then, only as time might permit the examination of alternative assumptions.

There are also two indirect ways in which the user can influence the evaluations. One way is by directly editing the assumptions that are used to explain contradictory data. Users can examine the assumptions associated with a given datum and add to or delete from the list of assumptions. In addition, users can directly change the belief levels associated with those assumptions.

The second indirect way to affect the results is to add entirely new possible kinds of evidence, or conclusions. Similarly, existing data items or conclusions can be deleted. This capability affects the results since any new data items or conclusions will impact the construction and evaluation of explanations.

All new entries and changes made by users are supported by straightforward interactions with the computer. The users are asked to answer relevant textual questions either by clicking on multiple choice answers or by typing text of their own. They do not need to understand SABER's internal representation schemes or be able to do any programming.

EXPLANATIONS TO THE USER

One of the reasons for exploring the EBR approach has been to try to obtain leverage from the internal generation of explanations in terms of the tool's ability to explain its reasoning to the user. This explanatory capability is evidenced at one level by the way that the internal explanations can be displayed to users. Thus, the internal structure of the explanations can be expanded into a textual and graphical representation giving the user a compact description of the explanations of strength, the confidence level associated with the most plausible explanation, all existing data, and a list of supports for each possible conclusion. The user can also see lists of assumptions that affect the evaluation of contradictory data.

At a higher level SABER is also able to give limited explanations of how it evaluated the possible explanations. Here, SABER invokes a separate explanation module which analyzes for the user how the various pieces of data in the current situation have been weighted to reach the current conclusion. This explanation includes the ability to give an indication of which data items may have changed in response to the user changes.

EXAMPLE

To make the above discussion more concrete, assume that SABER is set up to handle the data items and conclusions shown in figure 4. The data items are grouped to point and indicate the default conclusion. The decision problem in this case is to determine the probable intent of an aircraft that is of concern to the decision maker. The results of this determination of aircraft intent will ultimately be either to launch an attack, to take action to warn off the aircraft, to try to obtain further information, or to do nothing.

EXAMPLE DATA: GROUPED TO SHOW DEFAULT RELATIONSHIPS

CONCLUSIONS:	hostile routine	hostile threat	friendly	commercial
DATA ITEMS:	esm H	no radar	esm F	esm C
	radar H	no iff	radar F	radar C
	visual H	high speed	visual F	visual C
	IFF H	closing in	known F	air corridor C
		steep descent	IFF F	IFF C
				speed C

H = hostile F = friendly C = commercial

Figure 4. Example data that is grouped to show default relationships.

Next, assume that data items appear in discrete time steps, T1 through T5, in the following order: commercial air corridor, commercial speed, commercial IFF, closing in, steep descent. It is assumed that as these items enter the system, there is no other information available.

SABER's analysis at time T1 is shown in figure 5. The basic conclusion is that there is a slight preference for concluding that the aircraft poses a hostile threat. The underlying analysis is that the absence of expected data, such as, IFF and radar are important factors that push toward the hostile threat conclusion. The only contrary evidence, the commercial air path factor, is easily explained away by supposing that a hostile threat aircraft could use such paths as a subterfuge. This analysis shows the use of negative data items by SABER through which the tool can do some analysis based on the lack of some kinds of information (missing data items). It is thus, the absence of some kinds of data can be used in constructing an analysis of a given situation. It is also shown that the tool will reach some conclusion even when there is very little data to go on.

Ordered conclusions: hostile threat, commercial, hostile routine, friendly

Confidence: possible

EXPLANATIONS

1. hostile threat

DIRECT SUPPORT: no IFF, no radar

ASSUMPTIONS:

About air path C:
- hostile aircraft is imitating C
- hostile is off course

2. commercial

DIRECT SUPPORT: air corridor C

ASSUMPTIONS:

About no IFF:
- iff malfunctioning

About no radar:
- radar malfunctioning

3. hostile routine

DIRECT SUPPORT: none

ASSUMPTIONS:

About air path C:
- hostile aircraft is imitating C

About no IFF:
- hostile only testing our response

About no radar:
- hostile only testing our response

4. friendly

DIRECT SUPPORT: none

ASSUMPTIONS:

About air path C:
- friendly is off course

About no IFF:
- iff malfunctioning

About no radar:
- radar malfunctioning

Figure 5. Analysis at time T1.

The analysis at time T3 is shown in figure 6. The conclusion here is that the aircraft is commercial and is receiving a stronger preference. Basically, the analysis is now favoring the explanation that most completely accounts for the available data, in that, there are three pieces of data directly suggesting a commercial aircraft and only one points to a hostile threat.

Ordered conclusions: commercial, hostile threat, hostile routine, friendly

Confidence: possible

EXPLANATIONS

commercial 1.	hostile threat 2.
DIRECT SUPPORT: air path C, speed C, IFF C ASSUMPTIONS: About no radar: radar malfunctioning	DIRECT SUPPORT: no radar ASSUMPTIONS: About air path C: hostile aircraft is imitating C hostile is off course About speed C hostile aircraft is imitating C About IFF C hostile aircraft is imitating C

hostile routine 3.	friendly 4.
DIRECT SUPPORT: none ASSUMPTIONS: About air path C: hostile aircraft is imitating C About no radar: hostile only testing our response About speed C hostile training maneuver About IFF C hostile training maneuver	DIRECT SUPPORT: none ASSUMPTIONS: About air path C: friendly is off course About no radar: radar malfunctioning About speed C friendly training maneuver About IFF C friendly training maneuver

Figure 6. Analysis at time T3.

The analysis at time T5 is shown in figure 7. At this point in time, there are strong indicators that the aircraft is a hostile threat, and SABER's analysis points to that conclusion. The fact that an aircraft is closing in with a steep descent is taken to be much more significant than the other data items. That emphasis on the significance of those items may result from an initial weight assignment, or from a user directing the tool to reach this conclusion based on these items.

Ordered conclusions: hostile threat, hostile routine, commercial, friendly

Confidence: possible

EXPLANATIONS

hostile threat 1.

DIRECT SUPPORT: no radar, closing in, steep descent

ASSUMPTIONS:
- About air path C:
  - hostile aircraft is imitating C
  - hostile is off course
- About speed C
  - hostile aircraft is imitating C
- About IFF C
  - hostile aircraft is imitating C

hostile routine 2.

DIRECT SUPPORT: none

ASSUMPTIONS:
- About air path C:
  - hostile aircraft is imitating C
- About no radar:
  - hostile only testing our response
- About speed C
  - hostile training maneuver
- About IFF C
  - hostile training maneuver
- About closing in:
  - hostile testing our response
- About steep descent:
  - hostile testing our response

commercial 3.

DIRECT SUPPORT: air path C, speed C, IFF C

ASSUMPTIONS:
- About no radar:
  - radar malfunctioning
- About closing in:
  - C is off course
- About steep descent:
  - C is in distress

friendly 4.

DIRECT SUPPORT: none

ASSUMPTIONS:
- About air path C:
  - friendly is off course
- About no radar:
  - radar malfunctioning
- About speed C
  - friendly training maneuver
- About IFF C
  - friendly training maneuver
- About closing in:
  - friendly off course
- About steep descent:
  - friendly training maneuver

Figure 7. Analysis at time T5.

DISCUSSION

The SABER tool uses the EBR technique to model one strategy believed to be used by people in decision making. However, the SABER approach is not limited strictly to following the human model. Thus, SABER will examine explanations that can justify all of the possible conclusions, whereas people generally appear to consider only a few of the most likely seeming explanations. By analyzing all of the predefined conclusions, SABER avoids the problem of missing ways of explaining data, and it exploits the ability of the computer to construct the necessary explanations quickly. By presenting all possible explanations to the user, SABER will help overlook some of the explanations and thereby overcome some possible biasing problems. Of course, a limitation exists because SABER can only be used effectively in situations where there are a fixed number of known possible conclusions.

Even though SABER constructs explanations to support each possible conclusion, it is able to do so in linear time unlike the typical formal approach methods. The method of building up explanations to account for several pieces of data simply draws out the partial explanation from each piece of data that points to each possible conclusion, and combines those pieces. The weighting of each of the composite explanations is done when those pieces are combined so that the overall evaluation is done in linear time based on the number of pieces of data (figure 3).

Ease of use in training SABER to reach correct conclusions is promoted by the method of asking users to indicate how complete situations should be decided. The users do not deal at any point with specific weights for any data, or parts of data. When users are entering entirely new data, they are asked to indicate weights by means of the use of fuzzy terms rather than by supplying numeric values. Thus, the only expertise looked for in the user is related to the user's own knowledge of the decision-making scenario, not any knowledge of mathematical or computer theory. In addition, ease of use is being promoted by resisting addition of new features. Thus, we are continuing to find that some kinds of added functionality are needed, but our approach is to carefully weigh the added benefit of a new feature versus its cost in added user difficulties. It appears that too many computer tools start out as genuinely easy to use, but are then over-improved to the point that experts are needed just to use the tool correctly.

The overall use of SABER is dependent on the actual time constraints existing at any given time. At the most immediate level, the tool can always display the ordered conclusions and confidence levels. When more time is available, additional information can be obtained about underlying assumptions. The assumptions can then be manipulated to allow the user to determine what changes in the situation might lead to new preferred conclusions. Training and the entry of new data can be done in response to problems noted during actual use of the tool, but these actions are expected to take place at times when the tool is not in active use.

The question remains as to whether SABER can be shown to improve human performance. Improvement is looked for in SABER's ability to handle more information at once and will not be subject to human biases. It is also thought that by modeling a strategy that is intuitively familiar to users, the tool will be more readily understood and accepted by the users. The user will have the capability to modify the knowledge base to use his own terminology and to reflect changing circumstances. This should result in an improved ability to use the tool more easily and to better judge the quality of the tool's output. While work on exact interfaces is still in progress, we think the tool will assist users in keeping track of contradictory data to be more accurate in forming their own explanations of events. It needs to be demonstrated, however, that humans will benefit from using a tool like SABER. It is planned that SABER will be tested in real world Navy scenarios by using a simulation facility that can reproduce such scenarios in a laboratory setting.

CONCLUSION

The ability of SABER users to easily revise the knowledge base is seen as an advance in this kind of tool. That ability acts as an adjunct to the capability to train the tool and as a way of postponing obsolescence.

The capability to train the tool and revise the knowledge base is also seen as important factors in increasing user confidence in the tool. It is expected that users who have actively engaged in training the tool will have reason to be confident in the tool's analysis, and at the same time will have an appreciation for the limitations of that analysis. Confidence is further enhanced through the two-level explanation capability.

To some extent, these training features will lead to the tool being tailored to the individual abilities and built-in biases of the individual users. We think, however, that result is desirable as long as the user is actually an expert in the decision-making area. Thus, the point of trying to support naturalistic decision making is to help users do a better job of using their own decision-making strategies. It can only be expected that the tool would have to be somewhat individualized to accomplish that goal.

In general, we believe that human decision making can be improved through the use of tools like SABER, where the emphasis is on modeling the strategies actually used by people as a means of being able to directly influence when and how such strategies are used. We also believe that usability, and especially modifiability, by people who are not computer experts is extremely critical. If the actual users are able to have a direct impact on the results generated by the tool, they have a full appreciation for the tool's capabilities and only that kind of understanding can lead to successful use of the tool.

In summary, the SABER tool is being developed as part of an approach to improving human decision-making in time-constrained situations. The tool offers benefits directly related to the fact that the tool models one of the strategies of decision-making believed to be used by people. This approach should result in the tool being more easily understood and therefore more easily used and modified. The approach also lends itself to an ability to present information in a way that will be useful in overcoming possible cognitive biases. Thus, alternative explanations are always available and can be presented to users, along with underlying assumptions, in ways that can positively influence the decision making process.

REFERENCES

Charniak, E. 1991. "Bayesian Networks Without Tears," *AI Magazine*, 50–63.

Hammond, K. R., R. M. Hamm, J. Grassia, and T. Pearson. 1987. "Direct Comparison of the Efficacy of Intuitive and Analytical Cognition in Expert Judgment," *IEEE Transaction on Systems, Man, and Cybernetics*, 753–770.

Hirst, R. A., 1988. Uncertainty in Knowledge-Based Command Systems, Master's Thesis, University of Manchester, UK.

Kahneman, D., P. Slovic, and A. Tversky, 1982 (eds). *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press.

Klein, G. A., 1989. "Recognition-Primed Decisions," in *Advances in Man Machine Systems Research* 5, pp. 557–92, Rouse W. R. (ed.), J. A. Press.

Kolodner, J. L., 1991. "Improving Human Decision Making Through Case-Based Decision Aiding," pp. 52–68. *AI Magazine*.

Mitchell, T. M., R. M. Keller, and S. T. Kedar-Cabelli, 1986. "Explanation-Based Generalization: A Unifying View," *Machine Learning*, 1, 47–80.

Norman, D. A., 1988. *The Design of Everyday Things*, Doubleday/Currency.

Pennington, N. and R. Hastie, 1988. "Explanation-Based Decision Making: The Effects of Memory Structure on Judgment," pp. 521–533, *Journal of Experimental Psychology: Learning, Memory, and Cognition 14*.

Perrow, C., 1984. *Normal Accidents*, Basic Books.

Pratt, I., 1987. "Explanatory Asymmetry," *Proceedings of Alvey KBS Club Explanation SIG 3rd Workshop*, 168–182, University of Surrey.

Shafer, G., 1979. *A Mathematical Theory of Evidence*, Princeton University Press.

Smith, J. B. and M. Lansman, 1992. "Designing Theory-Based Systems: A Case Study," *CHI'92 Conference Proceedings*, 479–488, Monterey, California.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	December 1992	Final

4. TITLE AND SUBTITLE

DEVELOPMENT OF COMPUTER SUPPORT FOR NATURALISTIC DECISION-MAKING

5. FUNDING NUMBERS

PN: RM33D60/RM33T40
PE: 0602233N
AN: DN300090

6. AUTHOR(S)

D. C. Hair and K. Pickslay

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Naval Command, Control and Ocean Surveillance Center (NCCOSC)
RDT&E Division
San Diego, CA 92152–5001

8. PERFORMING ORGANIZATION REPORT NUMBER

TR 1558

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Office of Chief of Naval Research
Code 222
800 North Quincy Street
Arlington, VA 22217–5000

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

Technology Transfer Code M Military Application

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

In Navy command and control systems there are increasing problems with situations in which Navy tactical decision makers must make quick decisions based on uncertain information. The SABER tool is being developed to use an explanation-based reasoning approach to assist in making such decisions. Explanation-based reasoning is used to model an explanation-based decision-making process through which people are believed to perform some decision-making tasks. SABER constructs and evaluates alternative explanations to account for input data that may be incomplete and inconsistent. The explanations are evaluated according to three criteria: simplicity, completeness, and significance. SABER is designed to allow users to easily change the data and to modify the rankings of the explanations. This emphasis on modifiability effectively enables users to train SABER. The tool is being used to demonstrate to the Navy the capability of using AI techniques to assist humans in making critical command and control decisions.

14. SUBJECT TERMS

Cognitive Modeling
Decision Making
Decision Support
Displays
Man-machine Interface
Performance Measurement
Stress
Team Decision Making

15. NUMBER OF PAGES

25

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	SAME AS REPORT

21a. NAME OF RESPONSIBLE INDIVIDUAL	21b. TELEPHONE (Include A Code)	21c. OFFICE SYMBOL
D. C. Hair	(619) 553–5302	Code 444

INITIAL DISTRIBUTION

Code 0012	Patent Counsel	(1)
Code 144	V. Ware	(1)
Code 40	R. C. Kolb	(1)
Code 44	J. D. Grossman	(5)
Code 444	D. Smith	(1)
Code 444	R. Larson	(1)
Code 444	C. Hair	(5)
Code 961	Archive/Stock	(6)
Code 964B	Library	(2)

Defense Technical Information Center
Alexandria, VA 22304-6145 (4)

NCCOSC Washington Liaison Office
Washington, DC 20363-5100

Center for Naval Analyses
Alexandria, VA 22302-0268

Navy Acquisition, Research & Development
Information Center (NARDIC)
Washington, DC 20360-5000

GIDEP Operations Center
Corona, CA 91718-8000

Tactical Training Group Atlantic
Virginia Beach, VA 23461-5596

Tactical Training Group Pacific
San Diego, CA 92147-5080

Bureau of Navy Personnel
Washington, DC 20370

Office of Naval Research
Arlington, VA 22217-5000 (2)

American Psychological Association
Washington, DC 20002

Naval Sea Systems Command
AEGIS Program Manager
Washington, DC 20362-5160

Combat Systems Training Group Atlantic
Norfolk, VA 23521

AEGIS Training Center
Dahlgren, VA 22448-5190 (2)

U.S. Army Research Institute
Fort Bliss, TX 79906-0057

Armstrong Laboratory
Williams AFB, AZ 85240-6457