AD-A270 617

...ITATION PAGE

Form Approved
OMB No. 0704-0188

Publi... maint... suggi... 2220... ...hour per response, including the time for reviewing instructions, searching existing data sources, gathering and ...ation. Send comments regarding this burden estimate or any other aspect of this collection of information, including ...irectorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA ...Project (0704-0188), Washington, DC 20503.

1. AG... ...DATE: October 1993

3. REPORT TYPE AND DATES COVERED: Professional Paper

4. TITLE AND SUBTITLE: EXPLANATION-BASED REASONING IN DECISION SUPPORT SYSTEMS

5. FUNDING NUMBERS:
PR: CE16
PE: 0602233N
WU: 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

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES):
Office of Naval Technology
Office of Chief of Naval Research
Arlington, VA 22217

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

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 human technicians must make quick decisions based on uncertain information. The SABER tool is being developed using an explanation-based reasoning approach to assist in making such decisions. Explanation-based reasoning is modeled after 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 or inconsistent. The explanations are evaluated according to three criteria: simplicity, completeness, and significance. Users are given all of the explanations in ranked order. 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.

10P93

Published in *Proceedings of the 9th Annual Conference on Command and Control Decision Aids*.

14. SUBJECT TERMS:
Cognitive Modeling, Decision Making, Decision Support, Displays, Man-Machine Interface, Performance, Stress, Team Decision Making

15. NUMBER OF PAGES

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

93 10 8 0 24

UNCLASSIFIED

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

NSN 7540-01-280-5500

Standard form 298 (BACK)

UNCLASSIFIED

Explanation-Based Reasoning in Decision Support Systems

D. Charles Hair and Kent Pickslay
NRaD
Code 444
271 Catalina Boulevard
San Diego, CA 92152
email: hair@popeye.nosc.mil
phone: 619-553-5302
fax: 619-553-4149

ABSTRACT

In Navy command and control systems there are increasing problems with situations in which human technicians must make quick decisions based on uncertain information. The SABER tool is being developed using an explanation-based reasoning approach to assist in making such decisions. Explanation-based reasoning is modeled after 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 or inconsistent. The explanations are evaluated according to three criteria: simplicity, completeness, and significance. Users are given all of the explanations in ranked order. 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.

Explanation-Based Reasoning in Decision Support Systems

D. Charles Hair and Kent Pickslay
NRaD
Code 444
271 Catalina Boulevard
San Diego, CA 92152
email: hair@popeye.nosc.mil
phone: 619-553-5302
fax: 619-553-4149

INTRODUCTION

The Navy is faced today with increasing problems involving decision making based on incomplete and uncertain information. These situations often involve severe time constraints. Information can come in too quickly and in too high a volume to allow for full assimilation of the implications of the information.

Problems can arise from failures to react quickly enough, or from failures to correctly judge the significance of incoming information. Another problem is that human decision makers can be subject to undesirable biases which can lead them to misinterpret or misuse available data.

A computer tool is being developed at NRaD to assist human decision makers in such situations. The tool is named SABER (Situation Assessment By Explanation-based Reasoning). The SABER work is being done as part of the TADMUS (TActical Decision Making Under Stress) project, funded by the Office of Naval Technology.

The general aims of the TADMUS project are to explore the actual decision making strategies used by people in stressful situations, to determine precise ways in which those strategies fail, and to try to find ways to prevent the failures. As part of that effort, work is being done to develop new computer tools to be used in decision support systems.

The SABER work is focused on producing one of the computer tools to be part of the decision support system. It is important to note that SABER is not intended to replace the human decision makers. Another important goal of the SABER work is to develop a tool that can assist in reasoning with uncertainty in real time situations. As part of accomplishing that goal, the tool is expected to process data at a much faster rate than can be done by humans. In addition, a major emphasis has been put on developing a tool that can have its database and actual decisions easily modified by technicians who may not be experts either with computers or mathematics.

BACKGROUND

In situations heavily involved with modern technology it has often been the case that when things go wrong the blame is placed on human error. However, some recent research has suggested that in many cases the fault lies more in the design of the technological systems than in the system users (Perrow, 1984). The idea that humans in general do a poor job of making decisions, particularly in stressful situations, has also recently been called into question

(Hammond, 1987). The picture emerging from the TADMUS research is one in which people are generally believed to do a good job of decision making using their own innate decision making strategies. Experts in particular appear to have a high level of performance. What is suggested is that to try to improve that level of performance, what is needed is an approach that focuses on the actual naturalistic decision making strategies people use, rather than on other more analytic strategies.

There are three basic approaches in trying to use computers to assist human decision makers. One is to have the computer use methods that do not claim to be related to ways in which humans actually solve problems themselves, but which do claim to have formal validity. A second 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 as a means of influencing human decisions. The third approach is to model the way the computer reasons after the way humans are thought to reason.

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, 1979). While use of these methods has been quite fruitful, there are two problems that bear directly on the goals of the SABER work. First, those 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 those approaches have made real strides in making their resulting tools easy to use, it does not appear that a true layman could hope to 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, given sufficient time.

The second approach is basically a man-machine interface approach. An advantage of this approach lies in the fact that generally there is little extra computation time since the approach is only looking at ways to manipulate the interface. A drawback of this kind of approach is that the full potential for computers to be of assistance is not going to be realized if no use is made of computer techniques for modelling cognitive processes. Of course, much of the work involved with man-machine interfaces takes a hybrid approach by building some kinds of machine intelligence into the interface.

The SABER work emphasizes the modelling approach which is seen to have three possible benefits: (1) the computer may be able to use these methods more quickly to account for more data than a human can do; (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 approach is likely to be more easily understood by users, and that understandability should lead to greater user confidence as well as greater ease of user interaction and user modifiability. The biggest 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.

The modelling approach can produce tools that operate faster than the formal tools, although with a resulting decrease in the degree of optimality of the conclusions that are reached. This approach necessarily makes more use of the potential reasoning capabilities of the computer than the strict man-machine interface approach.

In addition to the emphasis on cognitive modelling, the SABER work is involved in an effort to explore naturalistic decision making processes as opposed to analytic processes. The two major naturalistic strategies that have emerged from various research in the TADMUS project are recognition-primed decision making (RPD) (Klein, 1989), and a story based strategy originally explored by Pennington and Hastie (1988).

Both of those strategies are being modelled in decision support tools being developed as part of the TADMUS project. The RPD tool is being developed at an outside laboratory, while the SABER tool is using the story based approach.

THE SABER PROJECT

The SABER work has focused on an idea referred to as explanation-based reasoning (EBR). This idea was first suggested by Pratt (1987), and later incorporated into a computer program by Hirst (1988). The initial impetus for the EBR work was to develop a new computer technique for reasoning with uncertainty. The work at NRaD has stressed development of the tool as one that models one of the known cognitive strategies used by humans in decision making. This focus is expected to yield a tool that is intuitively understood by users. It is also expected that by virtue of modelling a human decision making strategy, SABER is in a good position to be able to present information in ways that may be able to improve human performance by overcoming biases.

The EBR approach begins by assembling available data into explanatory structures. There is one such structure for each of the possible conclusions. Each of the explanations attempts to explain how every piece of data can be accounted for in support of each conclusion, even though some of the data items would naturally contradict reaching some conclusions. Contradictory data is explained through the use of internal assumptions. It is assumed that there are a fixed number of predefined possible conclusions, and that each data item points directly to one of those possible conclusions.

Once the explanations are constructed, SABER evaluates them to determine which seem the most plausible. Plausibility is based on three primary criteria: simplicity, completeness, and data importance. Simplicity is concerned with how well an explanation accounts for a conclusion without using extra assumptions. Completeness looks at how much of the available data is directly accounted for by the explanation. Figure 1 illustrates a simple case in which simplicity and completeness could come into play.

The relative importance of different kinds of data is considered in SABER through the use of a weighting system. 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. Where contradictory data are involved, weights assigned to the different assumptions used to explain contradictions are used to test the plausibility of applying those assumptions. Where 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.

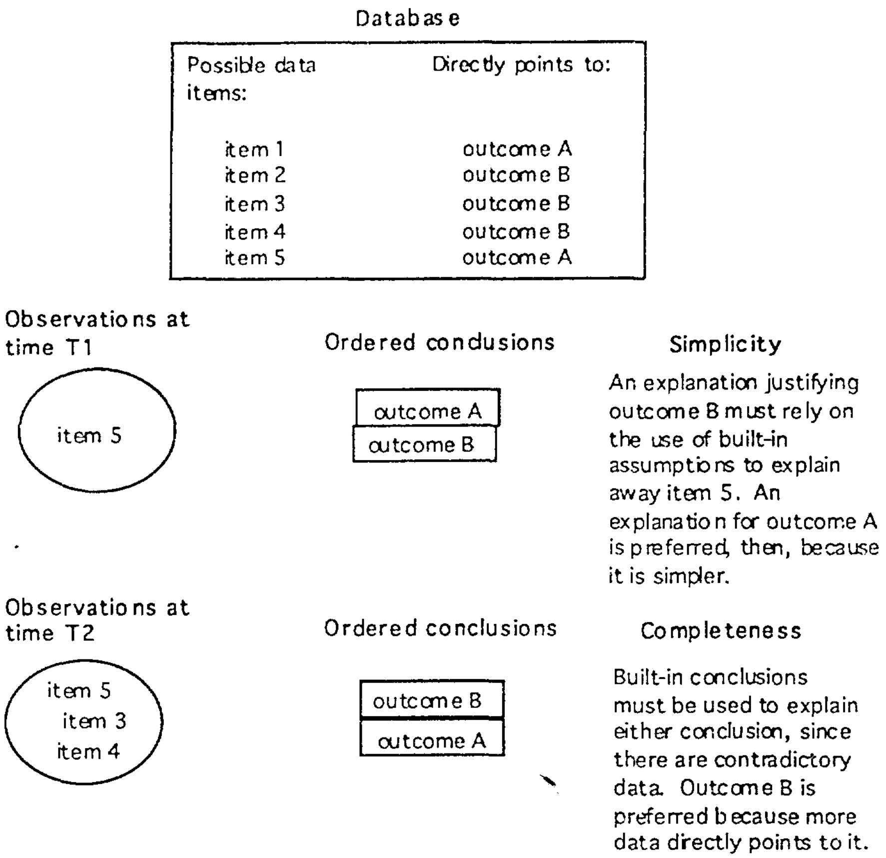

Figure 1. Explanation evaluation criteria.

This approach is quite similar to the explanation based decision making, or story based, strategy described by Pennington and Hastie (1988). They have developed data supporting the belief that in some situations people use a strategy in which they first put data together into a few story structures which explain the data. Decisions are then made based on a determination of which story appears to be the most plausible. There appears to be a strong similarity between the story structures suggested by that research and the EBR explanatory structures.

CONSTRUCTION OF EXPLANATIONS

As part of the internal knowledge base, SABER must know what the set of possible conclusions is. Thus, Figure 1 illustrates a situation in which there are only two possible conclusions, A and B. In addition, the knowledge base contains representations of each data item, where data items are represented as small sets of explanations.

Figure 2 illustrates a simplified data item representation. Here, the possible conclusions are that an aircraft is either friendly, neutral, or hostile. The figure shows that the representation of the friendly radar signal data item consists of a name, a weight, and three internal explanations. Each data item is expected to point directly to one of the conclusions, as is the case here where friendly radar signals obviously suggest the aircraft is friendly. So each data item has one internal explanation which is referred to as the default explanation. In addition, each data item must contain explanations that can be used to explain how each of the other possible conclusions could be reached despite the presence of this kind of data. These explanations take the form of sets of assumptions. So Figure 2 gives some simplified assumptions that could explain how a friendly radar signal could be associated with an aircraft that is really hostile or neutral.

The weight shown in Figure 2 is only included for illustrative purposes. The actual weighting system is more complicated, as it takes into account the relative importance of different kinds of data and the relevance of some kinds of data not being present. In addition, the assumptions are individually weighted so that users can indicate their degree of belief in given assumptions, and have SABER then reassess the plausibility of explanations based on those assumptions.

> 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: Typical Representation of a Type of Data

Figure 3 illustrates a simplified example of how the individual explanations in each data item representation are composed into larger explanatory structures when actual data comes into SABER. Here it is again assumed that there are only two possible conclusions, and the situation shows only two pieces of data have been received. The figure shows the composition of the explanation pieces from the data representations, but does not show which of the resulting explanations is the most plausible. In this case, there would be little for SABER to use in discriminating between the two explanations since the data items and assumptions are about equally divided. The plausibility determination would be based here on three factors. First, there would be consideration of which of the two data items is more important. Second, there could be an effect caused by the absence of other data items. Third, the assumptions themselves are weighted to have an impact on the plausibility calculation.

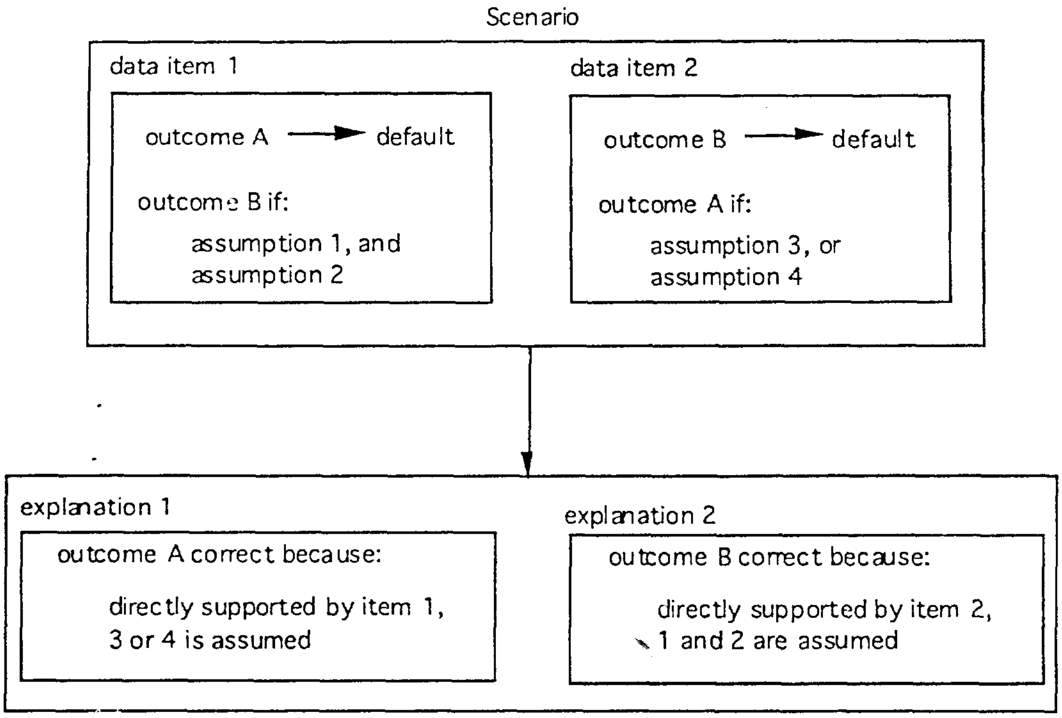

Figure 3. Schematic showing generation of combined explanations.

OTHER ASPECTS OF SABER

Among the features that we have incorporated into the EBR framework are the capability of being easily trained to reach new conclusions and handle new kinds of data, and the related capability of having the internal knowledge base modified through changes in the data item representations. We have also added a second level of explanation to try to provide extra assistance to users in understanding the tool's reasoning.

Users are able to directly modify both the database and the evaluation of explanations made by SABER. For the most part, user modifications are intended to be made off-line, at times when SABER is not in active use. 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.

The primary means users have to change the evaluation of conclusions is through the option of specifying which conclusion should be given the highest ranking for a given set of possible input items. SABER will then let the user know how the weights given to the data items can be changed in order to achieve the desired result. In this way users are allowed to effectively train the tool to reach correct conclusions. Thus, in the example shown in Figure 3 it would be possible during a training period for an expert user to specify which of the explanations should be treated as the more plausible.

Users are not asked to directly specify any weight changes, and are not expected to know the inner workings of the weighting system. The user simply specifies which conclusion to favor and which class of weights to change. Over a period of time, the overall set of weights are 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 over time the weights are expected to reach a near optimal setting.

In part, we believe that this tool will be easy to use and modify because of the fact that it is based on a naturalistic decision making strategy. That factor should work to make the tool's capabilities more understandable than is typically the case with more formal approaches. In addition, it is anticipated that the process of having users train the tool will increase their ability to understand both the advantages and limitations of the tool.

We are particularly concerned to have users understand the tool's limitations, because such understanding is a key both to establishing user confidence in the tool and to having users not believe that the tool is in any sense infallible. Thus, we want users to understand that they ultimately need to use their own decision making processes. This aspect represents another reason we believe that our modelling approach may work better than the formal approaches, since it is our experience that users of more formal tools are not inclined to question the results produced by such tools.

The goal of developing a tool that can process data quickly has been realized in the current version of SABER. The construction of explanations and the plausibility calculations are done in linear time based on the number of data items. Some ideas for future work may detract from speed of operation. For instance, some thought is being given to the use of case based reasoning and learning.

CONCLUSION

This work is leading to the production of a new computer decision support tool that models a decision making strategy used by human decision makers. The tool is expected to assist decision makers by presenting an explanatory overview of how given pieces of data can be related to each other, and by bringing out some of the assumptions underlying explanations that involve uncertain and contradictory data. One of the keys to successful use of the SABER tool lies in the ease with which laymen are expected to be able to modify the entire database and the conclusions reached by the tool. That ability to change and train the tool will in turn lead to increased user confidence in the tool. The work focuses on a fundamentally different idea than that found in most computerized decision support systems, because here the idea is to use the computer in ways that are expected to stimulate the natural decision making strategies of humans rather than to implement strategies that are seen to be improvements on the natural strategies.

REFERENCES

Charniak, E., 1991. Bayesian Networks without Tears, *AI Magazine*, 50-63.

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

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

Klein, G.A., 1989. Recognition-Primed Decisions, in *Advances in Man Machine Systems Research 5*, Rouse, W.R. (ed.), J.A. Press, 47-92.

Kolodner, J.L., 1991. Improving Human Decision Making through Case-Based Decision Aiding, *AI Magazine*, 52-68.

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

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.