A Tool that Uses Human Factors to Simplify Model Building and Facilitate More Accurate Strategic Decisions

Oscar Kipersztok,
Mathematics & Computing Technology
Boeing Phantom Works
P.O.Box 3707, MC: 7L-44
Seattle, WA 98124
oscar.kipersztok@boeing.com

Abstract

This paper describes a tool to validate hypotheses in strategic decision making. The system builds on experimental evidence of human factors that lead to more accurate decisions. The paper shows how those factors are translated into specifications for a user interface that simplifies the capture of expert knowledge to create complex models of strategic domains of interest. The tool uses graphical probabilistic reasoning, combined with text classification methods, to expedite the search for relevant information in a large number of documents to help validate hypotheses.

1 INTRODUCTION

This paper describes a tool that facilitates strategic decision making, i.e., decisions that, if wrong, can have significantly negative or even catastrophic consequences. The questions addressed are what type of information decision makers need to know, and how they process such information to arrive at decisions? Are there processes that lead to better decisions? What specific quantities do decision makers need to be able to make decisions?

There is evidence from experimental psychology that suggest trends and behaviors that, on the average, result in more accurate decisions. This work summarizes those human factors and processes and investigates how those can be used for building systems or computational tools that facilitate better decision making.

There has been a recent surge in technology used to support decisions. Graphical probabilistic networks (Bayesian networks), for example, are used to model large complex domains from the aggregation of smaller, local, probabilistic dependencies between the variables of a domain [Pearl, 1988; Lauritzen and Spiegelhalter, 1988]. The methodology uses efficient probability-update algorithms that guarantee consistent and correct propagation of uncertainty when the models are queried. In Data Mining [Cabena, et. al., 1998], algorithms are used to search for information in very large databases to discover patterns and trends that can explain new phenomena. For example, for strategic decisions, classification algorithms [Langley, 1996] are contributing to the ability to retrieve pertinent information in support of critical decisions. In many instances, not having the right information available at the right time can have serious consequences. The large amounts of information that can be accessed today make it difficult for decision makers to search for specific information that may be needed for time-critical decisions.

In this paper, a prototype tool is presented that uses human factors found to lead to more accurate decisions. The system uses graphical probabilistic reasoning, combined with text classification methods, to facilitate the search through a large number of documents. It allows the capture of user's knowledge in a straight forward manner to create complex models of domains of interest. These are also referred to as "mental models" of the domain.

As the user builds a model, a graphical probabilistic network is automatically created that allows the user to raise hypotheses by making queries to the model. At the same time a text classifier is also created for search and retrieval from a large volume of documents. Once built, the user queries the model and the system responds by predicting the likelihood of a hypothesis. In prediction-driven mode, the classifier will search for relevant documents that can be used to substantiate or negate a given hypothesis. In evidence-driven mode, the evidence is retrieved ahead of formulating the hypothesis. In both cases the search is driven by the most likely hypotheses or the most likely evidence. The tool creates a summary report describing the formulated hypothesis and the evidence gathered to substantiate or refute it.

This paper shows that it is possible to embed in the GUI design features based on the human factors mentioned to guide the operational use of the tool to flow in a manner that facilitates more expedient and effective strategic decision making.

2 HUMAN FACTORS FROM EXPERIMENTAL PSYCHOLOGY

Evidence from experimental psychology [Oskamp, 1965; Goldberg, 1968; Heuer, 1999] shows that factoring more information into a decision process, increases the confidence of the decision maker without necessarily increasing the accuracy of the decision. People tend to believe that they use considerably more information in their decisions than they actually do [Summers et al., 1970; Slovic and Lichtenstein, 1971]. People compile vast amounts of information, over a life time of making decisions. Although it may appear that all that information is brought to bear in each new decision; actually, only a relatively small amount is relevant for each individual decision [Shepard, 1964].

Over time, people develop "mental models" of the world that are comprised of ideas and concepts learned, and the relationships between them [Heuer, 1999]. Such mental models become the "filter" through which people interpret the world around them. There is a continuous refinement of the mental model based on new experiences gained that reinforce certain relationships and weaken others.

People tend to utilize two distinct modalities for decision making. One modality is referred to as the "mosaic" model [Heuer, 1999] where pieces of information are accumulated, in no particular pattern, on the belief that a unified picture will, eventually, emerge revealing the solution to a problem query. The emphasis is in accumulating as much information on the subject of interest as possible.

The second modality is using the mental model to create a hypothesis for the query and then use the hypothesis to guide the retrieval of information. The search for information is driven by the model and the information retrieved is used to validate, refute, or reformulate the hypothesis. Experimental psychology provides evidence that the latter is the modality that tends to conduce to more accurate decisions [Elstein, et. al., 1978]. The process of scientific discovery, for example, follows the second modality where theories are created about unexplained phenomena and experiments are designed to validate or refute the theories. In turn, the data are also used to improve theories and the process continues where experiment and theory affect each other.

The described system uses these results to guide the user in the process of making decisions. Firstly, it allows the user to build a mental model of their domain of interest in a fairly simple and unrestricted manner. Secondly, as the user queries the model, a hypothesis is proposed, and the system automatically identifies the essential parameters that are relevant to it. The system also conducts a search to retrieve information for use in validating or negating the hypothesis. The gathered evidence is used to redefine the query and to propose a better and more likely hypothesis. This cycle of operation patterns itself after the human factor trends described in this section that can lead to better decisions.

3 A KNOWLEDGE ACQUISITION PERSPECTIVE FOR BUILDING COMPLEX MODELS

From a knowledge acquisition perspective, the translation of the aforementioned human factors into system design specifications must conform to proper practices of user interface design. Domain models that facilitate strategic decisions can become quite complex. How can such complexity be managed during the knowledge acquisition stage, was an important question that needed to be addressed.

Two basic principles are defined and applied to the user interface design to manage the building of complex models. The first principle is driven by the goal to create a system that will allow direct and iterative interaction with the user, and will not necessitate the presence of an intermediary knowledge engineer. The communication language of the system with the user is limited exclusively to the language of the user's domain of interest. All technical terminology is hidden from the user. The user defines the domain as a list of concepts with appropriate labels and defines relationships between them with a weight of causal belief. The model is captured, as a directed graph and hidden from the user.

The second principle is to allow the user to attach specific qualitative and quantitative information to nodes and links in the graph. The intent is to associate the same type of information uniformly across all nodes keeping the amount of added information to a minimum. The new information, although of the same type, contains different semantic meaning for each individual concept or relation.

4 DESCRIPTION OF THE SYSTEM AND ITS USE

4.1 WHAT STRATEGIC DECISION MAKERS WHANT TO KNOW

Based on documents of strategic decisions and discussions during meetings of strategic nature, the observation was postulated that a common set of prediction parameters that decision makers are interested are the occurrence of events and the emergence of trends. Moreover, for each predicted event or trend, decision makers are interested in knowing their likelihood of occurrence, their magnitude or impact, and the time when they expected to occur.

4.2 BUILDING THE "MENTAL MODEL"

Specifying the use of mental models in computing architectures of decision support system requires defining methodologies of how such models are to be built. The type of model suggested here is a cognitive causal model in the form of a directed graph made up of nodes representing concepts and links representing relationships. Concepts are ideas represented by descriptive labels or phrases that convey a specific meaning or thought that is relevant to the domain of interest. Such model can be built by making a list of concepts and identifying, for each concept, all other concepts that may affect it.

This type of model is referred to as an “unconstrained model” because there are no restrictions as to the number of concepts or directional relations that can be defined in the graph. Any number of cycles is allowed. Since not all relations have the same degree of belief, a “weight of causal belief” is introduced and attached to each parent-child relation. Such weight represents the degree of belief that if the parent is true, it will influence the child to become true as well. An arbitrary scale between [0, 1] is used to assign such weight. In order to make inference about proposed hypotheses and make predictions about the domain, the unconstrained model is converted into a graphical probabilistic network by reducing the directed graph into an *acyclic* graph.

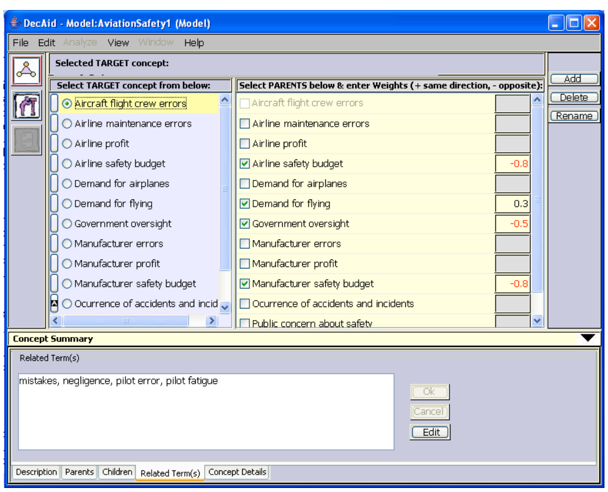

Figure 1 – Model building screen

During model creation the system allows as much flexibility for free association, without much concern about its ability perform inference. The use of the unconstrained model enables the system to acquire as much information as possible.

Figure 1, shows the model building screen for building the mental model. The model is built by defining concepts and their relations. As an example, concepts are shown for the “Aviation Safety” domain. In Figure 2, “Traffic Control Errors” and “Public concern” represent state transition concepts that affect the “Occurrence of Accidents”, a discrete event concept.

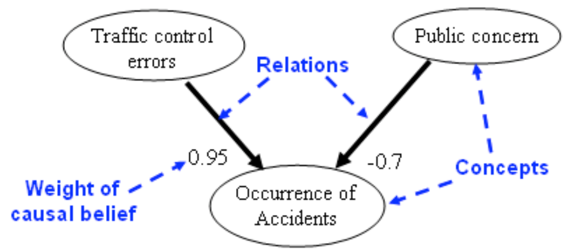

Figure 2- Examples of concepts, relations and weight of causal belief

The unconstrained model can be as large and complex as needed. The model is captured as a directed graph. Figure 3 shows a simple version of the Aviation Safety model where the negative signs at the end of the arrows representing negative weights of belief.

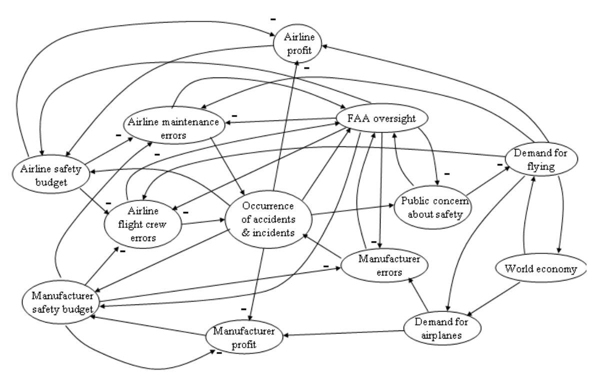

Figure 3- Simplified graph of Aviation Safety domain

4.3 INFERENCE AND PREDICTION

To be able to make predictions that decision makers are interested, such as the likelihood, magnitude and time of expected event and trend occurrences, Bayesian networks are used. The tool computes likelihoods and probabilities in response to queries. The user can add information to concepts to define them in term of dimensional quantities. Characteristic time quantities are also added for use in temporal reasoning [Nodelman, et. al., 2002]. To transform the directed graph into an acyclic Bayesian network, two steps are necessary. The first step is to eliminate cycles in the directed graph. This is done by minimizing information loss as a trade-off between the full expressiveness of the unconstrained model and the ability to make predictive inference. The second step is to utilize the weights of causal belief to create conditional probability tables between each node in the acyclic graph and its parent nodes. Similar approaches are found in [Rosen and Smith, 1996]

4.4 QUERY AND HYPOTHESIS CREATION

A query presents a question of how evidence or beliefs about current events (or trends) can predict the occurrence of future events of interest. The current evidence is the state of occurring events and on-going trends. Experimental psychology factors described in Section 2, suggests that only a few parameters are relevant in addressing a specific query in support of a decision. To answer a query the system identifies the sub model from the unconstrained model containing only those parameters that are relevant to the query [Druzdzel and Suermondt, 1994]. The screen in Figure 4 shows the results of such query with the corresponding prediction.

The response to the query is in the form of a hypothesis that predicts the likelihood of an event or trend given current evidence or beliefs. The user can simulate various ‘what if’ scenarios to gain insight into which of the postulated hypotheses will result in the most likely future events.

5 MENTAL MODEL GUIDING EVIDENCE SEARCH

As mentioned in Section 2, the iterative process of using mental models to raise hypotheses and the subsequent search for information to help validate them has been shown by experiments to result in more accurate decision making.

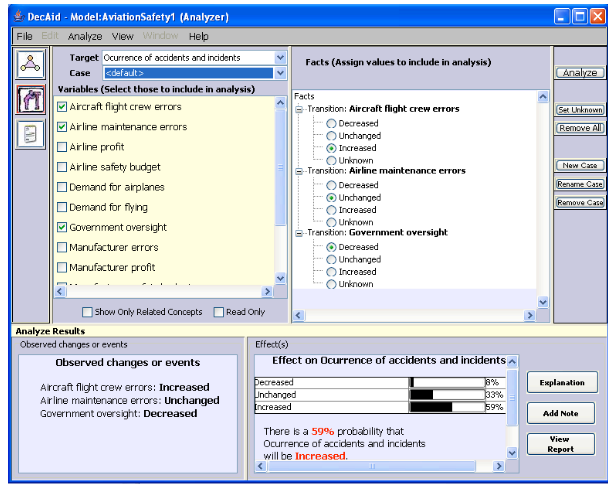

Figure 4- Query response with predicted hypothesis

While the user builds the model, the tool also provides added utilities to facilitate the automated building of a text classifier. The creation of the classifier is based mostly on the concept labels and the text descriptions that are attached to the nodes and links of the model graph.

The tool’s prediction capability using the mental model allows the user to discover the most likely hypotheses that justify the search for validating information. The capability to identify the parameters most relevant to the hypotheses further narrows the space of possible search. And the text classifier built using the domain language, introduced by the user, helps conduct the search in an efficient manner as is shown next. The search requires the presence of a large corpus of text documents. The text classifier built by the tool is used to rank text content relevant to concepts and relations associated with the hypotheses that resulted from the query.

Figure 5 shows the screen with the results of the classification process. Every document has been classified into the relevant concepts for the query. A score is computed by the classifier for each relevant concept. The documents are ranked according to content most relevant to the concepts in the hypothesis.

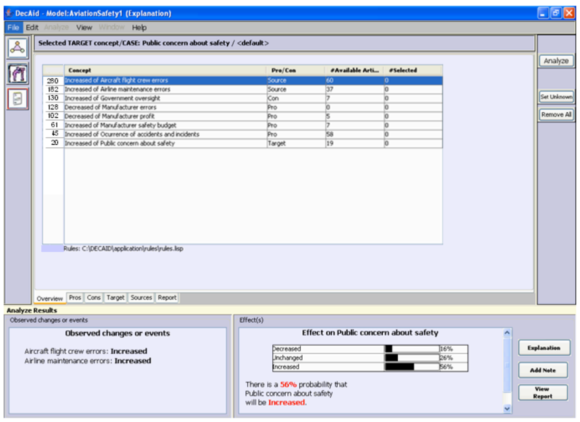

Figure 5- Query response with predicted hypothesis

The user can inspect each concept and identify the specific documents that have been associated with it. The screen in Figure 6 shows the specific documents that have been assigned to a concept. A second score is also computed to rank the documents with the most relevant content. Shown on the right pane of Figure 6 is information about author, source, title of each document, and a level of reliability that the user selects for each source. Such reliability factor is used to re compute the ranking of each document. Each document associated with a concept is further broken down into ranked paragraphs, as shown in Figure 7.

The highest ranked paragraph contains the content most closely associated with the concept. The user can, quickly, inspect the most relevant paragraphs of each document and select the paragraphs that contain evidence to help substantiate or reject the hypothesis. In this manner, the user can search through the documents and build a case by selecting the pertinent document paragraphs that help support or reject the hypothesis.

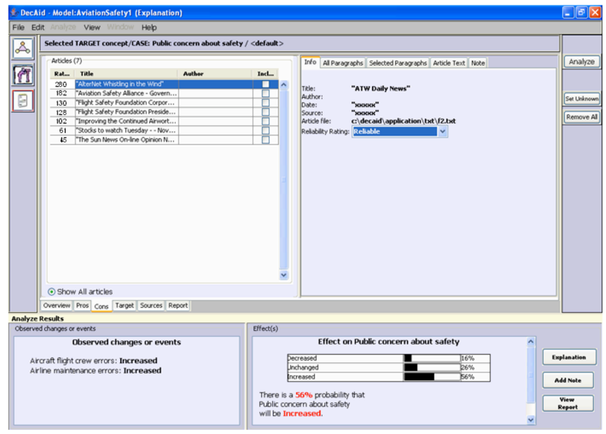

Figure 6 - Query response with predicted hypothesis

Figure 7 shows a screen with the ranked paragraphs for each specific document. The search process is driven, from the top down, by the most relevant content. Tens of thousands of documents can be processed in this manner, considerably reducing the search time.

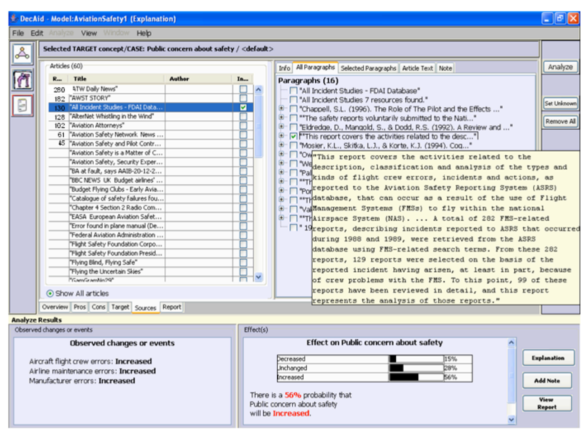

Figure 7 - Query response with predicted hypothesis

The last step is summarization of the evidence gathered from the selected documents and paragraphs which content helps to substantiate the hypotheses. The paragraphs selected by the user are compiled automatically and captured in a file with the postulated hypothesis and the conclusions drawn by the user.

6 SUMMARY

Experimental evidence from human factors studies describe behaviors that lead to more accurate decisions. A prototype system is presented that utilizes those factors in the design of the user-interface to help the user step through the creation of a model, presentation of a query, postulation of a hypothesis and, lastly, validation of the hypothesis by expediting the search through a large corpus of documents in the context of the most likely evidence.

References

[Allison, 1971] Graham Allison, Essence of Decision, 1971, *The Case Study Anthology*, Editor Robert K. Yin, Sage Publications, 2004.

[Cabena, et. al., 1998] Cabena P., P. Hadjinian, R. Stadler, J. Verhees, and A. Zanasi. *Discovering Data Mining: From Concept to Implementation*. Upper Saddle River, NJ: Prentice Hall, 1998.

[M. Druzdzel and H. Suermondt, 1994] Relevance in Probabilistic Models: "Backyards" in a "Small World". In *Working Notes of the AAAI 1994 Fall Symposium Series: Relevance*, pp 60-63, New Orleans, Louisiana, Nov 4-6, 1994.

[Elstein, et. al., 1978] Arthur S. Elstein et.al, Medical Problem Solving: An Analysis of Clinical Reasoning. Cambridge, MA and London: Harvard University Press, 1978, pp. 270 and 295.

[Goldberg, 1968] Lewis R. Goldberg, Simple Models or Simple Processes? Some Research on Clinical Judgments, *American Psychologist*, 23, 1968, pp. 261-265.

[Heuer, 1999] Richards J. Heuer, Jr., *Psychology of Intelligence Analysis.* Center for the Study of Intelligence, Central Intelligence Agency, 1999.

[Langley, 1996] Langley P. Elements of Machine Learning. San Francisco: Morgan Kaufmann, 1996.

[Lauritzen and Spiegelhalter, 1988] Lauritzen S. L., and D.J. Spiegelhalter. Local Computation with Probabilities in Graphical Structures and Their Applications to Expert Systems. *Journal of the Royal Statistical Society*, B, Vol. 50, No. 2, 1988.

[Nodelman, et. al., 2002] Nodelman U., Shelton, C. R., and Koller, D. Continuous Time Bayesian Networks. *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, pp. 378-387,

[Oskamp, 1965] Stuart Oskamp, Overconfidence in Case-Study Judgments, *Journal of Consulting Psychology*, 29, 1965, pp. 261-265.

[Pearl, 1988] Pearl. J. *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufmann, San Mateo, California, 1988.

[Rosen and Smith, 1996] Julie A. Rosen and Wayne L. Smith. *Influence Net Modeling with Causal Strengths: An Evolutionary Approach*, Proceedings of the Command and Control Research and Technology Symposium, Naval Post Graduate School, Monterey, CA, June 25-28, 1996.

[Shepard, 1964] R. N. Shepard, On Subjectively Optimum Selection Among Multiattribute Alternatives, in M. W. Shelly, II and G. L. Bryan, eds., *Human Judgments and Optimality,* New York: Wiley,1964, p. 166.

[Slovic and Lichtenstein, 1971] Paul Slovic and Sarah Lichtenstein, Comparison of Bayesian and Regression Approaches to the Study of Information Processing in Judgment, *Organizational Behavior and Human Performance*, 6 ,1971, p. 684.

[Summers et al., 1970] David A. Summers, J. Dale Taliaferro, and Donna J. Fletcher, Subjective vs. Objective

Description of Judgment Policy, *Psychonomic Science*, 18 , 1970, pp. 249-250.