# Requirements for Developing Strategic Decision Facilitation and Event Prediction Tools

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

## Abstract

This paper describes the requirements of a strategic decision facilitation tool that relies on forecasting to support critical decisions. A hypothesis-driven (data supported) system rather than a purely data-driven methodology. It further describes the importance of simple and natural human-computer interactions that simplify the creation of complex domain models in a system that uses probabilistic reasoning methods to facilitate high-quality decision making under uncertainty. Such a system helps users create complex models, query them for predictions, formulate hypotheses and validate their prediction with evidence retrieved from a corpus of text documents. The system must have a technology to automatically assemble and explain the forecasts so that users--who should not be required to understand the mathematics behind the forecast--will be able to understand why certain predictions are being made.

## 1 INTRODUCTION

A principal goal in any forecast of future events is to help decision-makers deal with uncertainty. Our pictures of the present and the past are always incomplete and noisy. So uncertainty is ubiquitous. A deeper concern is not just the future, but even our *theories and models* of how events will unfold that are underdetermined by our data. Thus, forecasting from data alone is not sufficient and it can be improved by embedding the forecast technology within a larger framework for decision support. Such a framework can supplement our raw data with information about the appropriate context in which to interpret the framework that allows for ongoing critical evaluation and validation of high-quality forecasts created.

forecast, will be able to better focus computational resources and minimize (as far as possible) the quantified uncertainty over the most relevant aspects of the forecast.

The goal of this paper is to clarify the requirements of a strategic decision facilitation tool that relies on forecasting to support critical decisions.

Such system will offer users maximum flexibility and provide quick turn-around through a decision facilitation process that allows: a) easy capture and organization of knowledge, b) building complex models that can be readily queried about future events, c) applying advanced algorithms, made transparent to users, to forecast predictions, d) searching and piecing together relevant and coherent argumentation in favor (or against) courses of action; and e) making actionable recommendations to facilitate significant strategic decisions.

## 2 KEY COMPONENTS

There are four key components to a forecasting system that will be discussed to facilitate high-quality decision-making: 1. the forecasting algorithms should have access to the context of decisions under consideration, not simply the raw data--that is, they should be *hypothesis-driven*; 2. the system should enable *simple and natural human-computer interactions* to allow forecasting directly over concepts of relevance and importance to the decision makers; 3. the simplicity of user interaction should not prevent the use of advanced *probabilistic reasoning methods* to quantify and minimize uncertainty over forecasts; and, lastly, 4. the system should be capable of *automatically constructing explanations* of forecasts which can be understood without requiring users to master the details of the forecasting algorithms. Together, these components yield a complete decision-support

# 3 HYPOTHESIS DRIVEN METHODOLOGY

## 3.1 HYPOTHESIS VERSUS DATA DRIVEN

There is a body of evidence in experimental psychology suggesting different modalities in the way people make decisions; some modalities result in more accurate decisions than others (Heuer, 1999). In general, there is the distinction between "data-driven" and "hypothesis-driven" decision making. In the former, the emphasis is on initial search and gathering of as much information as possible before raising a hypothesis leading to an informed decision. In the latter, the emphasis is on a more selective and guided information search driven by a prior hypothesis. An iterative process follows where the search is aimed at specific information enabling validation or rejection of the hypothesis. The hypothesis is either accepted with sufficient evidence or re- formulated based on insufficient evidence. Validated hypotheses with sufficient evidence, in general, lead to more accurate decisions. Our decision support system is architected to direct users to follow a process that in practice has been shown to result in more accurate decisions.

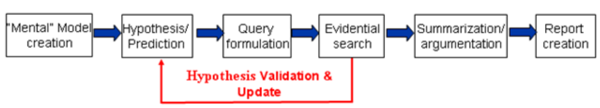


Figure 1: Sequential steps and feedback in hypothesis-driven decision making

# 4 SIMPLE AND NATURAL HUMAN COMPUTER INTERACTION

This section describes basic human computer interaction principles used to facilitate the creation of complex domain models while making transparent the complexity of analytic methods.

Figure 2 shows the three types of analytic methods required by the decision facilitation tool: 2a) knowledge representation and capture, 2b) reasoning inference and 2c) text processing and search methods.

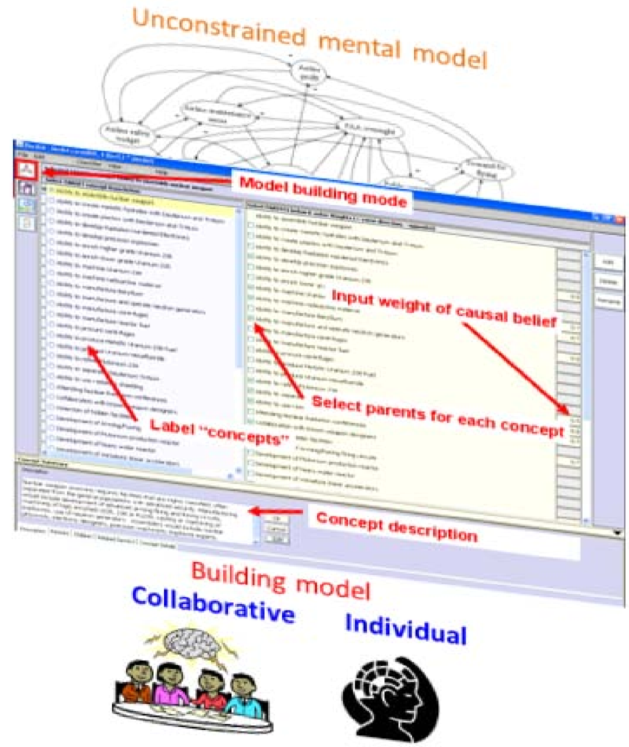


Figure 2a: Knowledge capture and representation

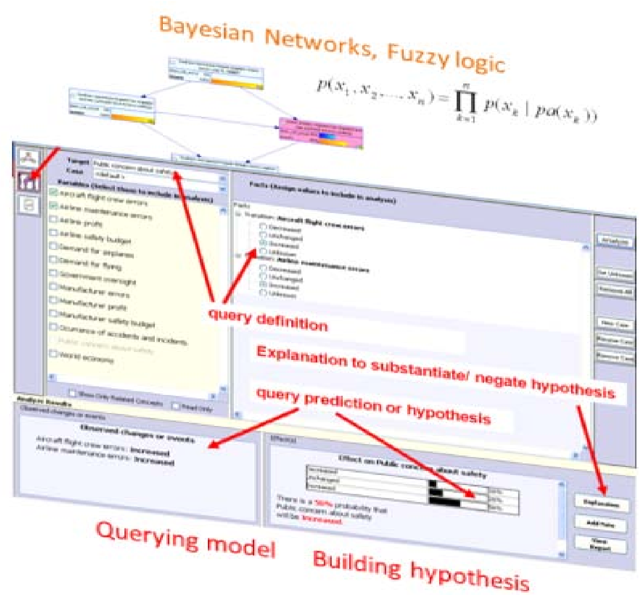


Figure 2b: Reasoning methods

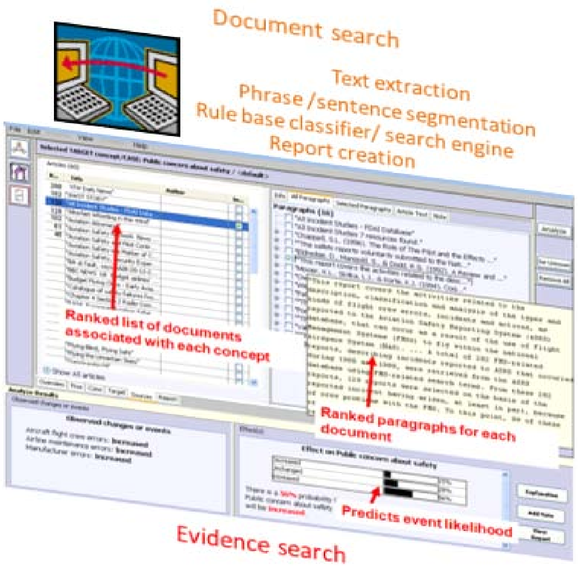


Figure 2c: NLP & Text Processing

Three corresponding screens have been designed as user interfaces to allow users to build models by defining concepts and their associations, allowing to query the model and make predictions by making assumptions based on existing facts or beliefs, and searching for information through a corpus of documents in order to validate the assumptions and predictions.

These methods are used to provide maximum flexibility and *ease of use* for *rapid* model creation, immediate query response and prediction, and fast document retrieval for forecast validation.

### 4.1 TRANSPARENCY OF ANALYTIC METHODOLOGY AND TERMINOLOGY

Advanced analytic methods often require familiarity with complex methodology. In our approach to modeling complex domains, it is not necessary for users to learn and familiarize with analytic methods.

By making the analytic methods transparent, users interact with the system by only using the familiar language of their domain. Domain concepts are defined using free language and users can add to those definitions to make their meaning more precise. This eliminates the need for knowledge engineers to acquire and convert user knowledge and expertise into computational models.

### 4.2 KNOWLEDGE AND RATIONALE CAPTURE

The model creation process is the most critical step. Users create a “mental model” of their domain, which consists of concept definitions and causal relations between them. Most concepts affect or are affected by other concepts. In most realistic domains there is feedback where a particular concept starts a causal chain feeding back to itself. Feedback loops can induce complex reinforcing or inhibiting dynamic behavior. The “mental model” is critical because it is used to make predictions and to process and interpret outside information.

### 4.3 FREE LANGUAGE, ASSOCIATION AND BRAINSTORMING

The use of free language in model creation enables more flexibility in building models. Concepts are defined and labeled with short phrases or using a few words. To reduce ambiguity, users further expand concept definition by providing added descriptions for more precise meaning. Concepts are defined based on specific assumptions that also need to be captured. Additional documents and information (e.g. names, locations, specific dates, events, etc.) are also associated with each concept for further clarification. The use of free language serves a dual purpose. Firstly, the words and phrases used to define the concepts are also used in the creation of a rule- based search engine to improve the recall and precision of retrieved content needed to substantiate the decisions. Secondly, the concept definitions and attached descriptions are also used to create chains of rationale that will provide explanations to subsequent predictions.

### 4.4 COMPLEXITY MANAGEMENT AND SCALABILITY

Automated knowledge capture should be made easy for the user. It should be just as easy to build models of high complexity as it is very simple models. The user should not be concerned with how the knowledge is being captured, represented and organized. Users should be able to add or subtract information to and from the model with ease and at will. Quantity of information should not be of concern to users. The information should be easily accessible at any time during the model building process, or later during the analysis phase. Providing flexibility to users during model creation in a free-associative, brainstorming fashion is important since it enables: a) adequate coverage of the domain, leaving no stone unturned; b) seamless scalability to large, complex domains; c) collaborative multiple-user participation with access to second opinions and feedback; d) ease of model refinement and evolution at any future time; and e) speed - quick addition and deletion of ideas without concern about performance or limits of scalability. Building models fast, with ease and with transparent complexity management, enables users to build unconstrained models of any size. The knowledge is represented using constructs that are readily mapped into graphical probabilistic networks for subsequent forecasting and predictive analysis.

# 5 PROBABILISTIC REASONING METHODS

This section describes the creation behind-the-scenes, following the construction of the unconstrained mental model, of a Bayesian network for making probabilistic forecasts of events and trends.

## 5.1 FORECASTING PROBABILITY, IMPACT AND TIMING OF EVENTS

Analysis and prediction methods must be compatible with knowledge representation and acquisition constructs used in the knowledge capture phase. In addition to defining concepts and relations, analysis methods require additional quantitative input parameters that need to be obtained from the user during model creation. In the spirit of devising easy ways to capture knowledge directly from users, the number of requested variables is kept to a minimum. Methods were developed to map those variables to fit the requirements of the analysis and inference algorithms. Quantitative inputs should be acquired from the user in an intuitive manner within the context of the familiar user's domain.

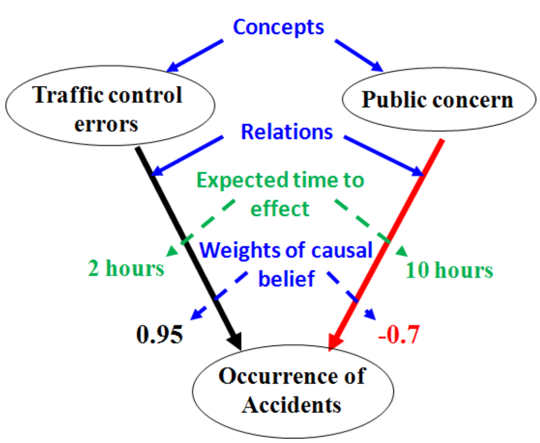


Figure 3: Request minimal number of parameters

The qualitative and quantitative inputs variables are shown in Figure 3. These are used to populate the parameters of graphical probabilistic (Bayesian) networks. The numbers represent the *weight of causal belief* that the user associates with each relation between pairs of concepts in the model. The relations and the belief numbers are used to build the structure and the conditional probability tables of the Bayesian networks by combining the weights of causal belief of incoming parent nodes (*Expected time to effect* will be discussed in section 5.2).

Alternative methods for building the networks and their conditional probability tables require obtaining probability estimates directly from domain experts for each combination of child and parent nodes' states. This tedious process can be aided by methods developed for probability elicitation (Wang, 2004). A major goal of our approach, however, is to circumvent this difficulty and make the process of building models readily accessible to users without need of expertise in graphical probabilistic methods. In either case, once the models are built, their performance and robustness can be validated using sensitivity analysis which can help identify the parameters that are most influential for any given query and prediction (Kipersztok and Wang, 2003).

Analysts and decision makers require probability estimates to guide strategic decisions. In order to maintain the simple and natural human-computer interface, our approach limits the decision space to predictions of *event occurrences* and *trends*. Decision makers need to know: a) how probable occurrences of events or emerging trends are; b) the magnitude of their *impact*; and c) the *time* when such events are expected to occur. Probabilistic models, in particular graphical models, provide capability to handle problems where data and information may be sparse, noisy or incomplete. In addition to rapid knowledge capture during model building, our methods also provide quick turn-around forecasting during the prediction phase.

As part of our on-going effort, we have built prototypes of the system. (Kipersztok, 2007) describes in more detail the implementation of various features of the system. The system has been applied to several specific domain areas. In (Seidler, et. al., 2010), the authors describe the DecAid system which was used to model and predict the readiness of a country to possess nuclear weapons capability. The paper reviews the domain associations used to build an unconstrained model. The model predictions were used to retrieve textual documents with information on Iran's nuclear program and to compile the risk assessment against the hypothesis that they are building a nuclear weapon.

## 5.2 REASONING ABOUT EVENT TIME

New methods are being developed that allow for probabilistic reasoning over systems evolving in continuous time (Nodelman, et. al., 2002). These techniques allow direct computation of distributions over when events of interest may occur. Moreover, they allow for automatic focusing of computational resources on those portions of the domain that may undergo rapid change. Larger, unified models over domains which include variables with widely divergent rates of change can thus be made computationally tractable. Machine learning algorithms can be used to help discover underlying structure in context where connections within the data are poorly understood (Nodelman, et. al., 2003).

There are already, in the literature, reviews of the advantages of these methods over traditional discrete-time probabilistic models--for instance, showing that the discrete-time models are subject to artifacts from the fixed time granularity and are less efficient to learn than the continuous time models (Nodelman, 2007). Adoption of these methods has been slow due to lack of exposure, limited software support, and ongoing research and development.

We define the *expected-time-to-effect* as one additional quantity to be provided by the user for each concept-pair relation defined in the model (Figure 3). Just as the *weight of causal belief* is used to create conditional probability tables in Bayesian networks; the addition of *expected-time-to-effect* is used for the construction of transition matrices in underlying continuous-time Bayesian networks. In this manner the system also can forecast the timing of events.

### 5.3 IDENTIFICATION OF CONCEPTS RELEVANT TO A QUERY

Experimental psychology experiments show that a critical number of variables is needed to make predictions at a fixed level of accuracy. Adding more variables to the decision increases the expert's confidence, without necessarily improving the accuracy of the prediction (Oskamp, 1965)(Shepard , 1964). This emphasizes the importance of being able to determine the critical, relevant concepts associated with a specific query. This important feature of our decision facilitation system is based on research done on relevance and feature selection learning algorithms (Druzdzel and Suermondt., 1994)(Fu and Desmarais, 2008b).

Querying the model triggers a prediction. A query is a request for predicting the future state of a 'target' concept given the assumptions about the current state of a set of 'source' or 'trigger' concepts. The model can contain any number of concepts and associations, but for each query the 'source', the 'target' and the set of 'relevant' concepts are the critical set of concepts that matter.

Once the model is complete the system is ready for inference and prediction, based on a specific query. At the time when the query is made, the system identifies the variables and relations relevant to that query and that subset of the unconstrained model is converted into a predictive model (Figure 4).

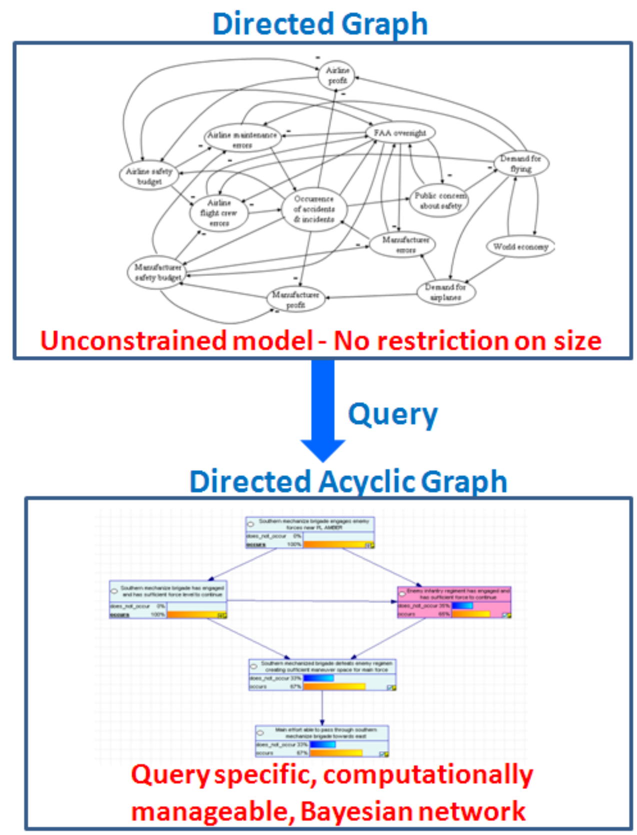


Figure 4: Creation of a predictive model (DAG)

### 5.4 INDIVIDUALS VS. COLLECTIVE JUDGMENT – CONSENSUS VS. DISENTING OPINIONS

A system that offers such model building flexibility and quick turn-around in decision-facilitation and forecasts can be equally effective for use by a single analyst as well as by a collective group of decision makers. Difficult and significant decisions are often arrived at by consensus in a group setting. Collective consensus is often built around a particular set of assumptions, a hypothesis, and a prediction. Once this occurs, it becomes increasingly difficult to deviate from the consensus opinion. Consensus can often be dominated by a vocal minority within a group at the risk of ignoring dissenting but equally, or more, valid alternative opinions.

In our system, various parallel hypotheses can be formulated with ease and subject to different sets of assumptions. With our proposed approach, a single team member may be capable of quickly making predictions and forecasting scenarios based on a dissenting hypothesis, while at the same time compiling evidence that can be used to steer the consensus opinion in a different direction that may, convincingly, lead to a different and possibly better decision.

# 6 FORECAST VALIDATION AND EXPLANATION

In (Lacave and Díez, 2002), the authors review various explanation methods for Bayesian networks and argue, on the one hand, that the *normative approach* for building expert systems, based on probabilistic reasoning, leads to more robust and accurate results. On the other hand, they also require more explanation capability because the methods are more foreign to human beings than in the *heuristic approach*.

Here, we are suggesting that much of the information to be presented as explanation during inference and prediction can be captured upfront, during the model building phase. It is part of the contextual knowledge imparted by the user as concepts and relations are defined. A very specific context and specific assumptions are made for every concept and causal relation defined in the unconstrained model. The *weight-of-causal- belief* and the *expected-time-to-effect* are quantities also defined subject to very specific assumptions. The system allows for user to systematically add such context during the model creation phase. The information is organized so that it is readily available for retrieval at the time that the explanation is needed.

## 6.1 EXPLANATION AND CHAIN OF REASONING

Predictions of highly probable events, of high impact and possibly occurring in the near future will be of most interest to users. Before courses of action are decided, decision makers require *explanations* that support a particular prediction and its assumptions. In addition, they also require convincing *evidence* that can back the predictions with plausible and believable facts. Qualitative explanations are provided by showing the causal chains of reasoning from trigger assumptions to predicted target outcomes where the entire relevant context that was captured during model building is organized and presented in every step along various paths of reasoning. Rationale captured and documented by the users, together with evidence retrieved from document search, constitute the basis for the explanation given to decision makers associated with forecasts and predictions by the system.

## 6.2 INFORMATION RETRIEVAL-VALIDATION THROUGH SEARCH

In political, cultural and socio- economic domains, validation often comes from validated evidential facts. Having a hypothesis makes the search more efficient because it narrows the search for specific information as evidence for clearly stated assumptions; thus, lending credibility and validity to the predictions. Precise concept definitions and rationale that explain concept relations together with gathered evidence from search make it possible to support a hypothesis-driven prediction. In arriving at critical decision, the facilitation methods discussed can help users step through a process that helps capture knowledge and data, organize them, invoke analysis methods to forecast predictions, piece together evidence, and rationale for or against courses of action, and make actionable recommendations. The final choice of action must be ultimately made by humans. The system will compute and present the necessary trade-offs between risk and cost for each recommended course of action.

Retroactive historical analysis constitutes another validation approach. It entails making predictions of past events and comparing the model forecast to actual outcomes. Predictions can also be compared among different methods.

## 6.3 RAPID PROTOTYPING - 'WHAT IF" SCENARIOS

Being able to quickly build complex mental models, and having the underlying machinery to automatically convert the created entities and relations into analytic models to make immediate predictions, provide single or multiple users with great flexibility. A single user can in one sitting use their knowledge to build a complex model, define concepts and relations, document their rationale, query the model to make predictions, and search for evidence to validate a new hypothesis. A quick turn-around decision facilitation method like ours enables users to postulate various 'what-if' scenarios and test parallel hypotheses side by side.

## 6.4 AUTOMATED DOCUMENTATION AND SUMMARIZATION

Our system automatically compiles and packages all the information needed for a strategic decision by summarizing the hypothesis and its assumptions, together with the associated evidence, the forecasts and the chain of reasoning explaining the prediction in the context of the specific concept definitions and the assumptions made when causal relations were defined. This capability to provide a summary documentation of the prediction, the assumptions and its explanation can be made available to second parties for critique and revision before actions of significant consequence are taken.

## 7 SUMMARY

Requirements for a decision facilitation system are presented that describe human- machine interface concepts that simplify for users the creation of complex domain models, while making transparent the analytic methodology that requires additional, specialized expertise. Those simplifying features are built into the user-interface to help users step through the creation of a model, query the model to make predictions, formulate hypotheses and validate the prediction from searched evidence (for or against) retrieved from a large corpus of documents. Explanation to predictions combines the rationale captured from the user during model development and the evidence gathered in support of a hypothesis; and it is presented to decision makers in context along the various paths of causal inference.

## Acknowledgments

Very special thanks to Uri Nodelman for his insightful comments and feedback; and for his valuable contribution in the area of temporal reasoning.

## References

Druzdzel, M., Suermondt, H. (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.

Fu, S. K. and Desmarais, M. C. (2008b). Tradeoff Analysis of Different Markov Blanket Local Learning Approaches. *In Proceedings of Pacific Asian Conference on KDD (PAKDD)*, 562-571, 2008b.

Heuer, R. J. Jr. (1999). Psychology of Intelligence Analysis. *Center for the Study of Intelligence, Central Intelligence Agenc*y, 1999.

Kipersztok, O. (2007). “A Tool that Uses Human Factors to Simplify Model Building and Facilitate More Accurate Strategic Decisions”. *Fourth Bayesian Modeling Applications Workshop at the Uncertainty in AI Conference*, Vancouver BC, Canada, July, 2007.

Kipersztok, O. and Wang, H. (2003). Validation of Diagnostic Models Using Graphical Belief Networks. *In Intelligent Systems for Information Processing: From Representation to Applications.* B. Bouchon-Meunier, L. Foulloy, R.R. Yager. Elsevier, 2003.

Koller, D., and Sahami, M. (2006). Toward Optimal Feature Selection. In *Proceedings of International Conference in Machine Learning (ICML)*, 284-292, 2006.

Lacave, C. and Díez, F. J. (2002). "A review of explanation methods for Bayesian networks", *Knowl. Eng. Rev.*, vol. 17, p.107 , 2002.

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

Nodelman, U., Shelton, C.R., and Koller, D. (2003). Learning Continuous Time Bayesian Networks. *Proceedings of the Nineteenth Conference on Uncertainty in Artificial Intelligence*, pp. 451-458, 2003.

Nodelman, U. (2007). Continuous Time Bayesian Networks. *Ph.D. Dissertation, Stanford University*. 2007.

Oskamp, S., (1965). Overconfidence in Case-Study Judgments, *Journal of Consulting Psychology*, 29, 1965, pp. 261-265, 1965.

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

Seidler, W., Kipersztok, O., and Wright, K. (2010). "Tools to Identify Nuclear Terrorist," *Journal of Radiation Effects, Research and Engineering*, Vol 28, No.1, pp 122 - 133, February 2010.

Wang H. (2004). Building Bayesian Networks: Elicitation, Evaluation, and Learning. *Ph.D. Dissertation, University of Pittsburgh*, 2004.