Simplifying Probability Elicitation and Uncertainty Modeling in Bayesian Networks

Patrick Paulson[*] and **Thomas E. Carroll** and **Chitra Sivaraman** and **Peter Neorr**
Stephen D. Unwin and **Shamina Hossain**
Pacific Northwest National Laboratory
Richland, WA

Abstract

In this paper we contribute two novel methods that simplify the demands of knowledge elicitation for particular types of Bayesian networks. The first method simplifies the task of experts providing conditional probabilities when the states that a random variable takes can be described by a fully ordered set. In this order, each state's definition is inclusive of the preceding state's definition. Knowledge for the state is then elicited as a conditional probability of the preceding state. The second method leverages the Dempster-Shafer theory of evidence to provide a way for the expert to express the degree of ignorance that they feel about the estimates being provided.

Introduction

Currently, system administrators must be intimately familiar with their cyber assets and their organization's missions. But as the network of cyber resources continues to grows, it becomes exceedingly difficult to adequately prioritize time and resources across possible threats as the crucial tie between cyber assets and organizational missions is absent from most cyber monitoring tools. As business needs and market pressures are causing cyber systems to become more interconnected and thus more susceptible to cyber attacks, organizations require a tool that allows them to gauge risk exposure from multiple *risk perspectives*, such as public safety, environmental impact, and shareholder return.

This need motivated us to develop Carim, a decision-support methodology that provides an assessment of the consequences of threats to components of cyber systems so that security personnel can better allocate resources to protect key components. Because of the evolving nature of cyber attacks, we've relied on non-probabilistic techniques to allow us to characterize the completeness of the knowledge used to make risk assessments.

Carim models each asset in a system as a particular asset type. Asset types have known mitigating relationships with other asset types. The mitigating relationships are elicited from domain experts and best practices and encompass a consensus view on the types of actions that can be taken to reduce an asset's vulnerability. For example, a workstation might have mitigating relationships that include the installation of anti-virus software, a backup server and related software, and a procedure for installing operating system patches. Each mitigating relationship involves other assets that might have mitigating relationships that require analysis. This network of mitigation relationships gives us a tool to elicit best practices from domain experts. It is similar to the causal mapping approach used for constructing Bayesian networks, where expert knowledge is represented as causal maps that are then, in turn, used to construct Bayesian networks (Nadkarni & Shenoy 2004). However, the elicitation of the conditional probabilities necessary for Bayesian networks proved difficult. This drove us to develop new methods for eliciting knowledge from experts.

Our Contributions In this paper we present two novel methods to simplify the demands of knowledge elicitation on certain types of Bayesian networks. The first method describes the values of a random variable using a fulled ordered set of states. In this order, each state's definition is inclusive of the preceding state's definition. The second method uses Dempster-Shafer theory of evidence to provide a way for experts to express uncertainty in the estimates being provided.

Related Work Domain experts are heavily relied on to provide information about probabilistic networks. Yet, these experts often struggle with the complexity of this responsibility. The problem of eliciting probabilities attracted the attention of many Bayesian network community (Druzdzel & van der Gaag 1995; van der Gaag *et al.* 1999; Olumuş & Erbaş 2004). The elicitation of the probabilistic values for reasoning under uncertainty is a critical obstacle (Druzdzel & van der Gaag 1995; van der Gaag *et al.* 1999; Olumuş & Erbaş 2004). Various methods have been designed to elicit probabilistic relations. But these methods tend to be very time consuming and are difficult to apply when many thousands of probabilities must be assessed.

While Bayesian networks have been used previously to reason about beliefs (see, for example, (Simon & Weber 2006; Simon, Weber, & Levrat 2007)), we generalize these methods and formally tie them to Dempster-Shafer (DS) theory of evidence. We simplify DS theory such that the focal elements of a node (i.e., the subsets with non-zero mass) are confined to the singleton plus a general "don't know." No other subset is assigned mass.

Organization The paper is structured as follow. In Sec-

[*] Corresponding author. E-mail: patrick.paulson@pnl.gov

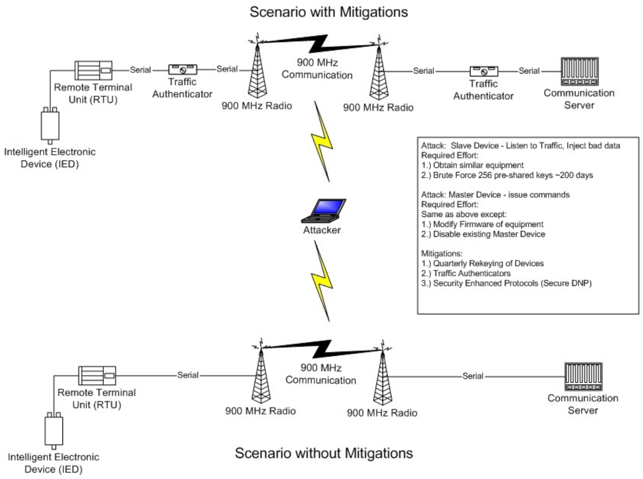

Attack: Slave Device - Listen to Traffic, Inject bad data
Required Effort:
1.) Obtain similar equipment
2.) Brute Force 256 pre-shared keys ~200 days

Attack: Master Device - issue commands
Required Effort:
Same as above except:
1.) Modify Firmware of equipment
2.) Disable existing Master Device

Mitigations:
1.) Quarterly Rekeying of Devices
2.) Traffic Authenticators
3.) Security Enhanced Protocols (Secure DNP)

Figure 1: A man-in-the middle attack scenario in which the attacker can eaves drop on the wireless communication channels.

tion we describe a scenario for which the Carim methodology was applied and discuss approaches in eliciting expert knowledge and representing uncertainty in the estimates that they provide. In Section we provide the background necessary to understand our contributions. We discuss our contributions in Section , which are novel methods for eliciting knowledge for Bayesian networks. Finally, we summarize and conclude in Section .

Eliciting SCADA Domain Knowledge for Carim

Carim has been applied to security in the domain of *supervisory control and data acquisition* (SCADA) networks, the networks used to control industrial processes. In this domain, we were particularly concerned with the possibility of a *man-in-the-middle* attack when a SCADA network included unsecured links between nodes. Figure 1 is a representation of this scenario. Our resident expert suggested two technical fixes that could be used, either independently or together, to protect against such an attack: SecureDNP, an encrypted wire protocol, and SSCP, a protocol that ensures data integrity. The effectiveness of these techniques depends on the "Rekey" policy used: how often the encryption and authentication keys are changed. Finally, the vulnerability to a man-in-the-middle attack depends on the capabilities of the attacker: an insider might have access to the required keys, and a state-backed attacker may have access to enough computing power to break the encryption scheme. The factors considered by Carim in assessing vulnerability to this attack are summarized in Figure 2.

In order to assess the vulnerability of the SCADA network using traditional Bayesian techniques, we would be required to elicit from our expert conditional probabilities for each combination of mitigation states and attacker expertise. In approaching our expert with this task, we realized that the expert was much more comfortable providing some values than others—for example, the relative effectiveness of rekeying policies was unknown, but the traffic authentication were clearer. When encoding these values into a Bayesian network however, the uncertainty of our expert disappears. While there is a good deal of controversy on the subject, the traditional approach of handling this problem with Bayesian networks is to ensure that the elicited probabilities encompass the doubts of the expert, and to not support additional "second order probabilities" (Pearl 1988, p. 360–363). We desired, however, to explicitly model uncertainty so that the end-user would have a measure of the applicability of the results. The method described here allows the expert to express uncertainty without forcing them to further analyze the factors causing their uncertainty so they can be expressed in one probability distribution.

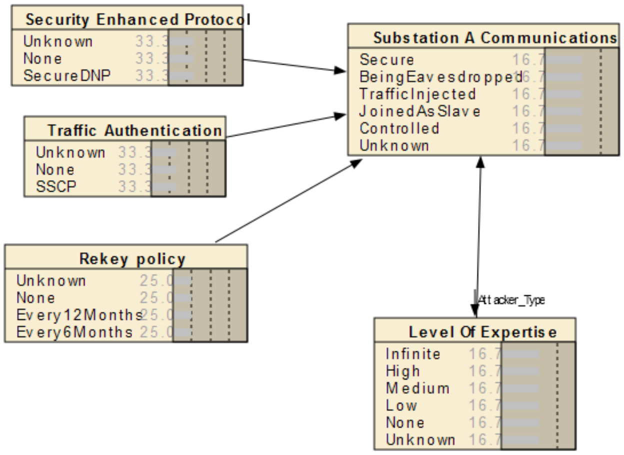

Figure 2: The elements for assessing the vulnerability of substation communications to an attack.

We also realized that we were putting undue burden on the expert by requiring them to state probabilities that had to match the constraints of the problem: it is always required, for example, that the vulnerability not decrease when the only change is that the expertise of the attacker increases.

We are charged then, with the following requirements:

1. Devise procedures to simplify the elicitation of probabilities that are constrained by additional factors
2. Model the uncertainty of the knowledge used to provide a measure of applicability of the model

A Technique to Simplify Elicitation

In order to apply a Bayesian network to this problem, the expert was required to provide *conditional probabilities* for each state of compromise of the asset for each combination

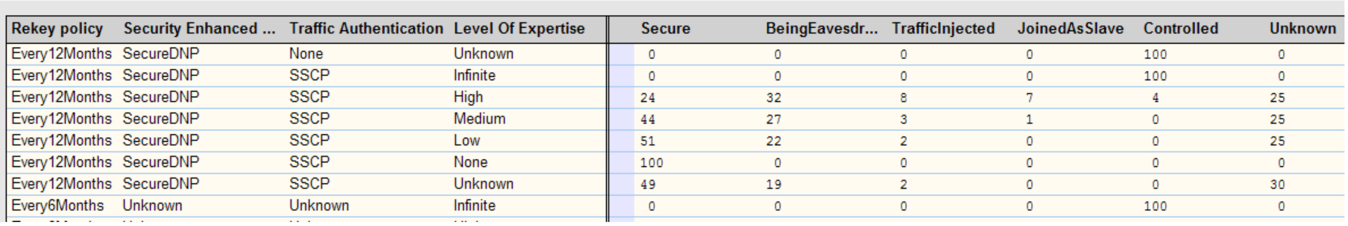

Rekey policy	Security Enhanced ...	Traffic Authentication	Level Of Expertise	Secure	BeingEavesdr...	TrafficInjected	JoinedAsSlave	Controlled	Unknown
Every12Months	SecureDNP	None	Unknown	0	0	0	0	100	0
Every12Months	SecureDNP	SSCP	Infinite	0	0	0	0	100	0
Every12Months	SecureDNP	SSCP	High	24	32	8	7	4	25
Every12Months	SecureDNP	SSCP	Medium	44	27	3	1	0	25
Every12Months	SecureDNP	SSCP	Low	51	22	2	0	0	25
Every12Months	SecureDNP	SSCP	None	100	0	0	0	0	0
Every12Months	SecureDNP	SSCP	Unknown	49	19	2	0	0	30
Every6Months	Unknown	Unknown	Infinite	0	0	0	0	100	0

Figure 3: Elicitation requirements for man-in-the-middle attack on substation communications. The left pane is the state space; the right is the sample space of the variable.

Table 1: The reduced elicitation requirements for substation communication security. These values are for the case when keys are changed every 12 months, both SecureDNP and SSCP are enabled, and the attacker has medium expertise.

Secure	1
Eavesdrop	60%
Inject	10%
Join	50%
Control	10%
Unknown	0.25

of the random variables that can affect the asset's state, as illustrated in Figure 3.

As described above, the expert is also allowed to specify a probability for the special state *Unknown*, which is probability they do not feel comfortable assigning to any particular state.

As can be seen in the Figure 3, the elicited probabilities in this problem have some interesting characteristics because of additional constraints on the state spaces of the variable. In particular, it is assumed that some states of compromise are "more difficult" to achieve then others; attackers with "higher" levels of expertise are accorded more probability of moving the asset into the more difficult states.

Because of these considerations we simplified our elicitation technique. For given states of the values of the mitigations and a specific level of expertise, we first have the expert give a estimate of the "uncertainty" they have in assessing the hypothesized situation. They are then asked to give, for the given level of expertise l, an estimate of the percentage of attackers with expertise l that can achieve the *lowest* level of compromise c_0 on the asset. (Since the lowest level of compromise is "completely secure", this value is 100 percent). Then, for each succeeding level of compromise c_i, they are then asked to estimate what percentage of attackers with expertise l who can achieve level of compromise c_{i-1} can also achieve level of compromise c_i. Section describes how we then convert these elicited values to probabilities used in a Bayesian network.

Using Figure 3 as an example, we are eliciting values for when keys are changed every 12 months, and both SecureDNP (encryption) and SSCP (authentication) are used. Using our technique, we elicited the values given in Table 1 for the case when the attacker has medium expertise. The values are elicited as percentages of the potential attackers with the given level of expertise that can move an asset to a more compromised level given the state of mitigations. Since we assume that all such attackers can leave the asset in the "Secure" state, the first elicited value is the percentage of attackers that can change the state to "being eavesdropped". In our example, the expert asserts a value of 60 percent. The next value we elicit is the percentage of attackers who can change the state to "inject messages." An attacker who can effect this change also has the expertise to eavesdrop. The expert testifies that 10 percent of all attackers who can eavesdrop can also inject messages. We continue eliciting values in this fashion in the order that the states are specified. Finally, we ask the expert to quantify her confidence in the values she provided. If the expert feels the amount of information given in the constraints is sufficient to determine the elicited values, than the "unknown" value would be zero. If the expert feels that they have no basis for their judgments, then "unknown" would be be one. Viewed this way, the "unknown" value is the portion of information required to make a judgment that is missing.

Background

In the following we briefly describe the foundations, Bayesian networks and Dempster-Shafer theory of evidence, on which we build our contributions.

Bayesian Networks *Qualitative Bayesian Networks* (Halpern 2003, p. 135), as a special case of *discrete influence diagrams* (Kjaerulff & Madsen 2008, p. ix), are convenient to elicit and encode an expert's impressions of factors that influence values in their domain of expertise. In order to be operational, *quantitative* Bayesian network requires a myriad conditional probabilities to be specified for each combination of values in an expert that they may not feel comfortable in estimating.

A Bayesian network $N = (G, P)$ over a set of random variables $\mathcal{X} = \{X_1, \ldots, X_n\}$ consists of a directed acyclic graph (DAG) G that encodes a set of conditional independence assertions and local probability distributions P for each variable. Together, G and P form a joint probability distribution over $\mathcal{X}$.

To be a Bayesian network, N must possess the local Markov property. Denote by $\mathrm{pa}(X_i)$ and $\mathrm{nond}(x_i)$ the set of parents and non-descendants, respectively, of X_i. A network possess the *local Markov property* if, for each $X_i \in \mathcal{X}$,

$X_{\mathrm{pa}} \in \mathrm{pa}(X_i)$, and $X_{\mathrm{nond}} \in \mathrm{nond}(X_i)$, the proposition (Neapolitan 2004, p. 37)

$$P(x_i) = 0 \vee P(x_{\mathrm{pa}}|x_{\mathrm{nond}}) = 0 \vee \\ P(x_i|(x_{\mathrm{pa}}|X_{\mathrm{nond}})) = P(x_i)$$

evaluates to true. In words, the local Markov property states that each variable is conditionally independent of its non-descendants given its parent variables.

The local Markov property makes Bayesian networks an effective technique for eliciting knowledge: by viewing the network, an expert can determine if all factors are being considered when determining the probability of an event.

Dempster-Shafer Theory The inability to express uncertainty is a drawback of the approaches based on probability theory (Halpern 2003, p. 24). However, expressing uncertainty is a necessity when attempting to elicit understanding in knowledge-poor domains (see, for example, (Forester *et al.* 2004; Donell & Lehner 1993; O'Hagan & Oakley 2004)). In contrast to purely probabilistic methods for capturing domain knowledge, Dempster-Shafer theory (DS) provides a rich mechanism for describing the range of beliefs about a result (Gordon & Shortliffe 1990). This richness comes at the expense of complexity in both eliciting the values for expressing the different types of ignorance and in the combination of multiple pieces of evidence (Ai, Wang, & Wang 2008).

In the following we summarize DS theory. We refer the reader to (Gordon & Shortliffe 1990) for a reference on DS theory. Let X be a random variable specified by the finite set $\mathbf{X}$ of its values. Set $\mathbf{X}$ is also called the *frame of discernment*. A *basic probability assignment* (BPA) $m_{\mathbf{X}}$ over $\mathbf{X}$ is a function

$$m_{\mathbf{X}} \colon 2^{\mathbf{X}} \to [0, 1],$$

where $2^{\mathbf{X}}$ is the power set of $\mathbf{X}$, for which

$$m_{\mathbf{X}}(\emptyset) = 0 \qquad \text{and} \sum_{S \subseteq \mathbf{X}} m_{\mathbf{X}}(S) = 1.$$

The *mass* or *degree of belief* $m_{\mathbf{X}}(S)$ of S is a measure of that portion of belief that is committed exactly to S by $m_{\mathbf{X}}$ and not to any particular subset of S. Each subset S such that $m(S) > 0$ is called a *focal element*. There are two measures that bound the interval that $m(S)$ resides. The function

$$\mathrm{bel}_{\mathbf{X}}(S) = \sum_{T \subseteq S} m(T)$$

computes the *belief* (or *support*) for all $S \subseteq \mathbf{X}$. The *plausibility* of each $S \subseteq \mathbf{X}$ is given by

$$\mathrm{pl}_{\mathbf{X}}(S) = \sum_{T \cap S \neq \emptyset} m(T).$$

Belief measures the total mass that is constrained to move within the set of interest, while plausibility measures the total mass that can visit somewhere in the set of interest but can also move outside it. From the definitions, we see that $\mathrm{bel}_{\mathbf{X}}(S) \leq m_{\mathbf{X}}(S) \leq \mathrm{pl}_{\mathbf{X}}(S)$.

In the next section, we discuss our methods for eliciting knowledge from experts.

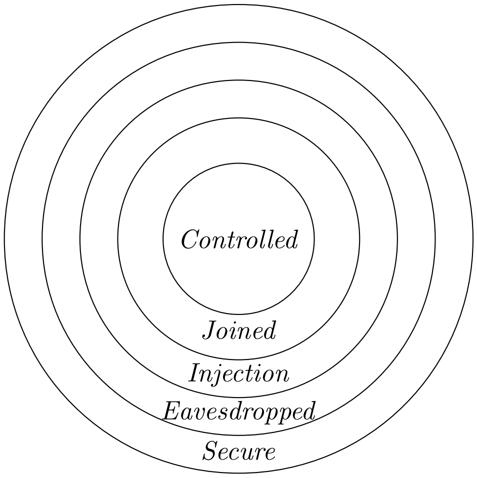

Figure 4: The inclusive compromise states of the substation communications using set theory.

Method

We next discuss our contributions to eliciting knowledge from experts. The first contribution is when the states of a variable can be described by a fully ordered set. In this set, a state implies all the preceding states. Our second contribution is using Dempster-Shafer theory of evidence to allow experts to express uncertainty of the estimates that they provide.

Simplifying Conditional Probability Elicitation in Bayesian Networks

Our goal is to elicit values in the form shown in Table 1 and calculate conditional probabilities that can be used in a Bayesian network. As an example, consider Figure 3 in which we need to elicit the probability of compromise conditioned on the Rekey policy, wire protocol, data integrity protocol, and the attacker's level of expertise.

The elements of the sample space of the random variables in a Carim model often belong to a simple order. For example, when considered in terms of the progression of an attack, the states of compromise in Figure 2 can be ordered as *Secure* $\prec$ *Eavesdropped* $\prec$ *Injection* $\prec$ *Joined* $\prec$ *Controlled*. What this says is that for the attacker to control devices, she must have joined the network, and to join, she must have the ability to inject traffic, and so on. The implication of states can be represented as inclusive sets as we have done in Figure 4.

An advantage of this constraint when eliciting knowledge is that we can state our elicitation in terms of an already elicited value, which eases the cognitive load on the subject matter expert. For example, instead of asking: "What is the likelihood that that a person with high skill level can eavesdrop on the network", and then separately asking "What is the likelihood that a person with high skill level can inject traffic into the network", we can ask "What is the likelihood that a person with high skill level who can eavesdrop the network can also inject traffic into it?"

Let $s_1, \ldots, s_n$ be states of X such that state s_{i+1} implies s_i (i.e., s_{i+1} is a subset of s_i). We elicit beliefs $P(s_1), P(s_2|s_1), \ldots, P(s_n|s_{n-1})$ from experts given the parents of X. But in probability theory, the elements of the sample space of a random variable must be disjoint. We obtain the disjoint sample space by defining x_i to mean for $s_i \wedge \neg s_{i+1}$. Treating the beliefs as probabilities, the probability $P(x_i)$ of X taking the value x_i given X's parents is:

$$P(x_i) = (1 - P(s_{i+1}|s_i))P(s_1) \prod_{j=2}^{i} P(s_j|s_{j-1}), \qquad (1)$$

for $i = 1, \ldots, n-1$, and

$$P(x_n) = P(s_1) \prod_{j=2}^{n} P(s_j|s_{j-1}). \qquad (2)$$

If X is conditionally dependent on other variables, we have all the necessary values to construct a Bayesian network to compute $P(x_i)$.

Implementing Subset of Dempster-Shafer Theory with Bayesian Networks

The greatest disadvantage of DS theory is that in contrast to probabilistic models, which are described by their respective density functions, DS models must be described by a set, which grows exponentially with the number of variable values. It would be difficult to elicit degree of belief for each and every set. If we can represent that problem with a graph that satisfies the Markov property, we then can use the computational efficiency of Bayesian networks to compute degrees of belief.

Beliefs are elicited from experts for each value x of variable X and also the element $Unknown$, which is equivalent in DS theory to the set $\mathbf{X}$. All other sets have no mass. Given these conditions, $P(x)$ satisfies the requirements of $m_{\mathbf{x}}$ as $\sum_{\bar{x} \in \mathbf{X} \cup \{Unknown\}} P(\bar{x}) = 1$. A Bayesian network can be constructed such that the node that represents X has a state for each of its focal elements. The node's conditional probability table comprises the elicited conditional probabilities of X given its parents. The network output for the node computes $P(\bar{x})$, for each $\bar{x} \in \mathbf{X} \cup \{Unknown\}$. The belief in x is simply $\text{bel}_{\mathbf{X}}(x) = P(x)$ and the plausibility is $\text{pl}_{\mathbf{X}}(x) = P(x) + P(Unknown)$.

We now consider an example. There are three variables X, Y, and Z, where X conditionally depends on Y and Z and Y and Z are conditionally independent. From the definition of joint probability, the probability $P(\bar{x})$ of $\mathbf{X} \cup \{Unknown\}$ is

$$P(\bar{x}) = \sum_{\substack{\bar{y} \in Y \cup \{Unknown\} \\ \bar{z} \in Z \cup \{Unknown\}}} P(\bar{x}|\bar{y}, \bar{z})P(\bar{y})P(\bar{z}).$$

This is the probability computed by the Bayesian network.

Conclusion

In eliciting knowledge for Carim, we frequently came upon the situation where we needed to determine the subjective probability of a member of a simple order according to a domain expert. For example, the probability that a threat will compromise an asset at a particular level of compromise. Additionally, the domain expert may have the ability to know when their knowledge about an area is incomplete, but be unable to further describe the characteristics of the incomplete knowledge. For these reasons, we wanted our users to be aware of the completeness of the knowledge in decisions.

We solved these problems by using Bayesian networks constructed using knowledge gained via our elicitation methods described in this paper. The first method simplifies the elicitation of conditional probabilities when the sample space of a random variable can be described by a fully ordered set of inclusive states. The conditional probability of a state is dependent only on is predecessor. The second method implements a subset of Dempster-Shafer theory using a Bayesian network. This allows the network to provide a measure of uncertainty along with its output.

References

Ai, L.; Wang, J.; and Wang, X. 2008. Multi-features fusion diagnosis of tremor based on artificial neural network and D-S evidence theory. *Signal Processing* 88(12):2927–2935.

Carley, K. M., and Palmquist, M. 1992. Extracting, representing and analyzing mental models. *Social Forces* 70(3):601–636.

Davey, B. A., and Priestley, H. A. 2002. *Introduction to Lattices and Order*. Cambridge University Press, 2 edition.

Donell, M. L., and Lehner, P. E. 1993. Uncertainty handling and cognitive biases in knowledge engineering. *Systems, Man and Cybernetics, IEEE Transactions on* 23(2):563–570.

Druzdzel, M., and van der Gaag, L. 1995. Elicitation of probabilities for belief networks: Combining qualitative and quantitative information. In *In Uncertainty in Artificial Intelligence (95): Proceedings of the 11th conference, Los Altos CA*, 141–148. Morgan Kaufmann.

Fiot, C.; Saptawati, G.; Laurent, A.; and Teisseire, M. 2008. Learning bayesian network structure from incomplete data without any assumption. In Haritsa, J.; Kotagiri, R.; and Pudi, V., eds., *Database Systems for Advanced Applications*, volume 4947 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg. 408–423. 10.1007/978-3-540-78568-2_30.

Forester, J.; Bley, D.; Cooper, S.; Lois, E.; Siu, N.; Kolaczkowski, A.; and Wreathall, J. 2004. Expert elicitation approach for performing atheana quantification. *Reliability Engineering & System Safety* 83(2):207–220.

Gordon, J., and Shortliffe, E. H. 1990. The Dempster-Shafer theory of evidence. In Shafer, G., and Pearl, J., eds., *Readings in Uncertain Reasoning*. San Mateo, California: Morgan Kaufmann Publishers. 529–539.

Halpern, J. Y. 2003. *Reasoning about Uncertainty*. Cambridge, Massachusetts: The MIT Press.

Heckerman, D. 1997. Bayesian networks for data mining. *Data Mining and Knowledge Discovery* 1:79–119.

Kjaerulff, U. B., and Madsen, A. L. 2008. *Bayesian Networks and Influence Diagrams: A Guide to Construction and Analysis*. Springer.

Nadkarni, S., and Shenoy, P. 2004. A causal mapping approach to constructing bayesian networks. *Decision Support Systems* 38:259–281.

Neapolitan, R. E. 2004. *Learning Bayesian Networks*. Upper Saddle River, NJ: Prentice Hall.

O'Hagan, A., and Oakley, J. E. 2004. Probability is perfect, but we can't elicit it perfectly. *Reliability Engineering & System Safety* 85(1–3):239–248.

Olumuş, H., and Erbaş, S. O. 2004. Determining the conditional probabilities in Bayesian networks. *Hacettepe Journal of Mathematics and Statistics* 33:69–76.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.

Simon, C., and Weber, P. 2006. Bayesian networks implementation of Dempster Shafer theory to model reliability uncertainity. In *Proc. of the 1st International Conference on Availability, Reliability and Security (ARES '06)*.

Simon, C.; Weber, P.; and Levrat, E. 2007. Bayesian networks and evidence theory to model complex systems reliability. *Jounal of Computers* 2(1):33–43.

van der Gaag, L.; Renooij, S.; Witteman, C.; Aleman, B.; and Taal, B. 1999. How to elicit many probabilities. In *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence*, 647–654. Morgan Kaufmann Publishers.