A Dynamic Bayesian Network for Diagnosing Nuclear Power Plant Accidents

Thomas B. Jones[1], **Michael C. Darling**[1,2], **Katrina M. Groth**[2],
Matthew R. Denman[2], **and George F. Luger**[1]
[1]Departments of Computer Science and Electrical & Computer Engineering
University of New Mexico, Albuquerque, New Mexico, USA
[2]Risk and Reliability Analysis, Sandia National Laboratories, Albuquerque, New Mexico, USA

Abstract

When a severe nuclear power plant accident occurs, plant operators rely on Severe Accident Management Guidelines (SAMGs). However, current SAMGs are limited in scope and depth. The plant operators must work to mitigate the accident with limited experience and guidance for the situation. The SMART (Safely Managing Accidental Reactor Transients) procedures framework aims to fill the need for detailed guidance by creating a comprehensive probabilistic model, using a Dynamic Bayesian Network, to aid in the diagnosis of the reactor's state. In this paper, we explore the viability of the proposed SMART proceedures approach by building a prototype Bayesian network that allows for the diagnosis of two types of accidents based on a comprehensive data set. We use Kullback-Leibler (K-L) divergence to gauge the relative importance of each of the plant's parameters. We compare accuracy and F-score measures across four different Bayesian networks: a baseline network that ignores observation variables, a network that ignores data from the observation variable with the highest K-L score, a network that ignores data from the variable with the lowest K-L score, and finally a network that includes all observation variable data. We conclude with an interpretation of these results for SMART procedures.

Introduction

The ability to model sequential data is important both in science and engineering. The Bayesian Belief Network (BBN) is an important generalization of the hidden Markov model, as it offers a factored form of the full Markov representation (Pearl 1988; 2010). Thus, besides offering a more expressive representation for the components of a complex causal system, it also supports tractable inference (Murphy 2002). Recently Dynamic Bayesian Networks (DBNs) have been used in continuous speech recognition, protein sequencing, bioinformatics, and more. We are applying this technology to the states of a nuclear power plant reactor.

Although there has been use of AI technology (Kim and Bartlett 1996) in analysis of the components of the nuclear power generation process, there has been no research to date

analyzing the use of the dynamic Bayesian network technology to monitor the full complex processes involved in nuclear power generation. So we begin this paper describing why we feel this technology is not only appropriate, but how it can add immense value to the nuclear power generation monitoring process.

First, the DBN may be composed of both probabilistic and knowledge-based processes. The conditioned aspect of such systems comes from both testing individual components, such as sensors, as well as from the ascribed relationships between components. The knowledge-based side comes from the original construction of the DBN, that also enables factored Bayesian network computational efficiencies, as well as from assigning specific probabilities to component relationships when these are known by human domain experts.

Second, the DBN system is able to run situation analyses in better-than-real-time, supporting both current event and alternate scenario tracking. Third, with the flexible time discretization process available to the DBN, the power generation system may be monitored intensely, by every second, say, when situations are critical, as well as more slowly, by minutes or hours, when the reactor state supports this.

Finally, the DBN is able to automate the complex (and often manual) current processes for analysis of critical accident scenarios (see discussion of Severe Accident Management Guidelines, next). This automatization process is very important as, once it is verified as optimum, it will allow the human reactor monitors to see in real time what the alternatives to particular situations are and what steps are recommended (rather than having to look these situations up in complex paper manuals as is the current protocol).

Pearl (2010) suggested this automated DBN scenario analysis approach when he proposed causal "counterfactual algorithms." For Pearl, "counterfactual" meant to test situations that are not "currently true" for a particular situation. We plan to use this logic to explore alternative scenarios, where, given a particular critical state, possible solution alternatives can be explored.

Though nuclear power plant accidents are extremely rare, the effects can be harmful for people, the environment, and the economy (FNI 2012). The steps taken in the event of an accident are critical for limiting the extent of the damage to the plant and its surrounding environment. Nuclear

power plant operators are given manuals with step-by-step *Emergency Operating Procedures* to follow for each type of anticipated accident. However, for unanticipated, beyond-design-basis accidents there are Severe Accident Management Guidelines (SAMGs) which are developed from expert judgments and best-estimates analyses. If the plants monitoring instruments fail, which was the case during 2011 Fukushima accident (FNI 2012), operators must act with severely limited information about the current state of the reactor.

In this paper, we present a prototype analysis tool for a sodium fast reactor using the SMART (Safely Managing Accidental Reactor Transients) procedures framework (Groth et al. 2014; 2015). This model would provide the ability to diagnose the likely state of the reactor and also provides insight into the relative diagnostic value of each of the plant's observed parameters. We explore the viability of the SMART procedures framework with cross-validation and the Kullback-Leibler divergence equation.

In order to gain an understanding of the possible responses of nuclear power plants during severe accidents, we simulated accidents with various conditions using the nuclear reactor accident simulation software SAS4a (Cahalan, Tentner, and Morris 1994). These simulations provide a rich data set that can be used for training a supervised machine learning algorithm. Using the simulated data, we built a dynamic Bayesian network in order to learn the conditions of observed power plant parameters that can lead to *Transient Over Power* and *Loss Of Flow* accidents. Transient Over Power is an accident sequence that involves an unintended increase in power. In a Loss of Flow accident the coolant circulation slows due to pump failure or an obstruction which results in the core overheating.

In the following section we describe how our model was developed and conditioned. We discuss the generation of data through simulation of accident sequences, the creation of the DBN, and demonstrate its interface. Then we evaluate the performance of our Bayesian model through analysis of the Kullback-Leibler divergence of its parameters, and analysis of the F-scores. Finally, we summarize our project and suggest future research.

Developing the DBN for Nuclear Reactors

Probabilistic Risk Assessments are used by the United States Nuclear Regulatory Commission to quantify the causes, likelihood, and consequences of nuclear accidents. *Dynamic Probabilistic Risk Assessment* studies a system's dynamics often by employing reactor state simulators. These simulators can provide detailed insight into specific accidents, but cannot yet be used to support real-time accident diagnosis.

The SMART Procedures framework involves creating a dynamic Bayesian network based on dynamic probabilistic assessment of simulated nuclear accident data to provide fast-running diagnostic support (Groth et al. 2014). This approach, as described in Figure 1, provides tools for diagnosing the likely state of a nuclear reactor given the values of the observed plant parameters. This will potentially allow for a greater understanding of the state of a reactor during accidents where only a subset of information might be available.

This understanding can enhance operators' decision making ability, especially during beyond-design-basis accidents.

Existing SAMGs rely on expert judgment and best-estimate analyses in order to capture the physical responses of the plant. However, procedure developers cannot anticipate every possible accident scenario. This limitation can be addressed with the use of *Dynamic Probabilistic Risk Assessments* coupled with *Discrete Dynamic Event Trees* to provide comprehensive coverage of the potential accident scenario space.

Dynamic Probabilistic Risk Assessment approaches can explore thousands of accident scenarios that form the basis for comprehensively learning the values of observable reactor parameters given known accident scenarios. This information provides a science-based support for operators especially during unanticipated accidents. These simulations also take into account the possible actions of plant operators from the plant's procedure manuals as parameters for the simulation.

Discrete Dynamic Event Tree based software dynamically branches the accident simulations whenever there is more than one possible outcome. The result of this approach is a comprehensive data set. However, the amount of data generated provides too much information to process during severe accidents. Therefore, as a preprocessing step, we discretized the data to allow for faster analysis.

Accident Data Generation and Processing

The data was generated for accident sequences using the SAS4a liquid metal reactor simulator which performs deterministic analysis on nuclear accident scenarios. We ran simulations with four target variables: the functional capacity of the nuclear power plant's main coolant system (differential pressure), Direct Reactor Auxiliary Cooling System (DRACS), the Balance of Plant (BOP) systems, and control rod insertion (SCRAM). These systems, seen in Figures 2, 3, and Table 1, were chosen as target variables since compounded failures of these system have the potential to lead to core damage.

The differential pressure parameter had three possible states: 100%, 50%, and 0%. This variable describes the capacity of the pumps removing heat from the core. The DRACS parameter had three possible states: *available*, *unavailable*, and *enhanced*. Balance of Plant has three states: *operational*, *decay*, and *shutdown*. The *SCRAM* parameter had three states: control rods *nominal*, *fully in*, or *fully out*. The control rods in a nuclear power plant are used to control the fission rate. The *SCRAM* parameter describes the level of the operators' attempt at an emergency core shutdown. The states of these target variables are initial conditions for the simulations.

The results of the SAS4a simulations is data which is comprised of 7189 distinct scenarios over a two hour simulation window each containing 2558 time steps. The simulations contain all possible combinations of target states that would lead to significant deviations in accident progression.
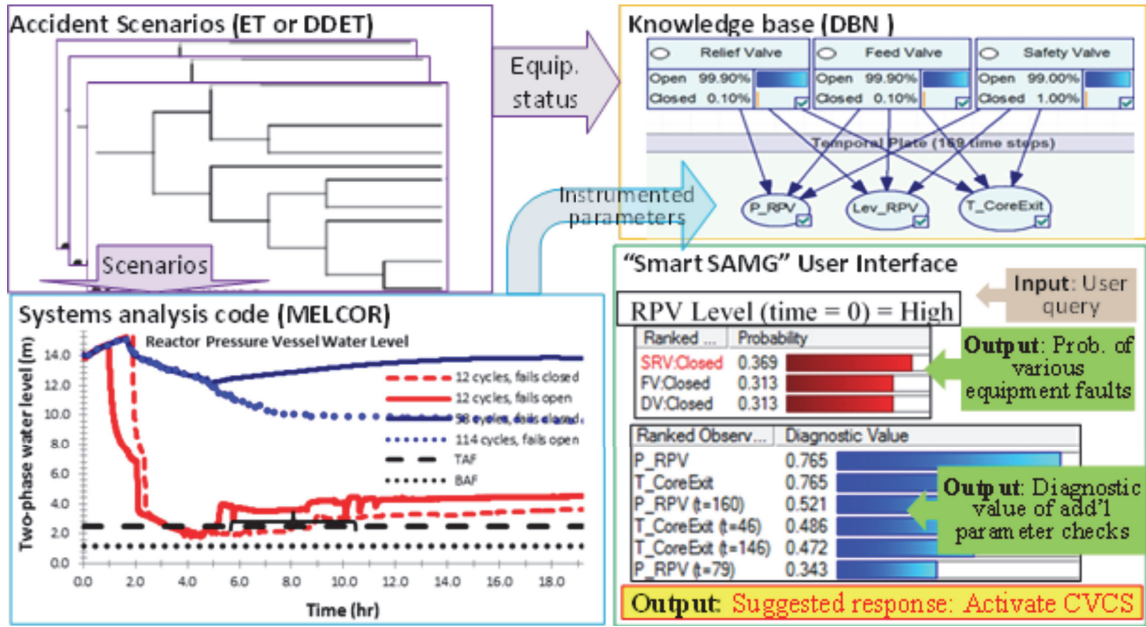
Figure 1: *Risk-informed "Smart SAMG" development process for nuclear power plant diagnostic support* Accident scenarios (upper left) are generated using event trees, the scenarios are simulated with a system analysis program such as MELCOR or SAS4a (bottom left), the data is used to generate a Bayesian network (upper right), the BN is analyzed using the GeNIe software program to diagnose the reactor's likely state (bottom right)

Discretization

The proper discretization of data values can improve the performance of supervised learning algorithms (Yang and Webb 2002). We discretized the conditional probability states of each variable into equal width distributions. This binning method calculates the maximum and minimum values for each variable's probability distribution and distributes the values into k bins of equal width. We divided the probability values for each variable into 3 bins.

Simulating nuclear power plant accidents with *Discrete Dynamic Event Trees* allows for comprehensive modeling of the nuclear accident space. However, due to the complexity of these models, they cannot be simulated and processed in real accident time. Therefore, the SMART procedures framework calls for the use of BNs to provide a means to utilize the data for faster-than-real-time decision support (Groth et al. 2014). Once built, the BN can be analyzed using the GeNIe (Druzdzel 1999) interface to inform the situational awareness of the operators (GeNie is a graphical interface to the University of Pittsburghs Structural Modeling, Inference, and Learning Engine or SMILE).

Construction of The Bayesian Network

Each simulation offered a permutation of the initial conditions. These initial conditions were drawn from the *Discrete Dynamic Event Trees* and were tied to a particular state of the target variables. The resulting states of the observation variables — variables that could be discovered by an operator's instrumental observations — are returned by the SAS4a simulator. Using this data we calculated the conditional proba-

bility of each of the observation variable states given each combination of possible target variable states at each time step.

With the conditional probabilities calculated, we then constructed the BN. Our system reads a provided dynamic BN outline which the user constructs with the GeNIe software. This pre-built BN contains a node for each of the plant's parameters, including the target and observation variables, and defines the relationships between them (see Figure 2). The observation variables, along with their states are given in Table 1. The system then populates the model with the conditional probabilities of each observation node at each time step (see figure 3).

Besides providing an interface for constructing the Bayesian network's structure, GeNIe also provides a platform to analyze the BN by propagating evidence and diagnosing the states of the plant's parameters. This functionality can be used as a decision support system. Users can input a set of known conditions, which would propagate evidence to the unobserved target variables. The posterior probability can be used to predict the evolution of important reactor systems.

Model Performance

In this section we study the effectiveness and performance of our DBN. In order to measure the pertinence of each of the plant parameters, we implemented a variable evaluator based on Kullback-Leibler (K-L) divergence (Cowell 2001). K-L divergence measures the distance between two probability distributions (in this case, between two BN models). In
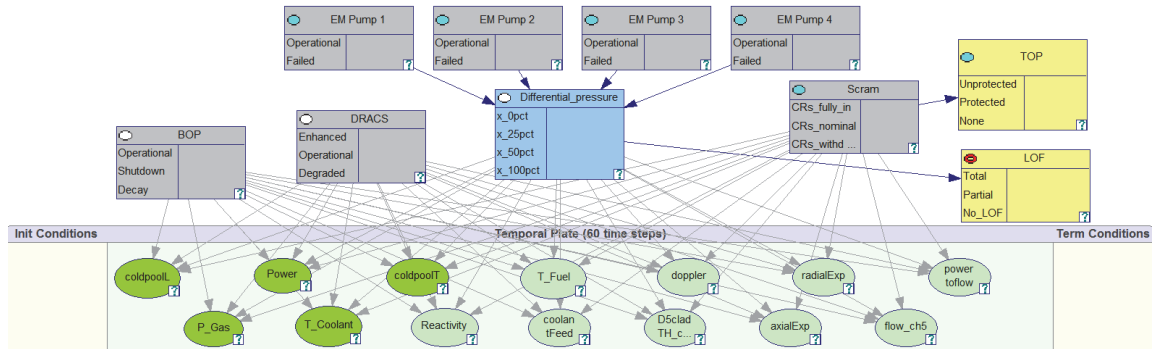
Figure 2: User-defined DBN Model Structure: the user creates the DBN model structure to define the relationships between observation nodes, target nodes, and intermediate nodes. The oval-shaped nodes, which are placed on a temporal plate, represent observed variables (the darker nodes are plant parameters observed in the control room, whereas lighter nodes are variables that are not directly monitored). The rectangular nodes represent reactor systems. The light rectangular nodes on the far left are targets representing the accident states. The system reads this hand-made model and populates its conditional probability tables based on the data.
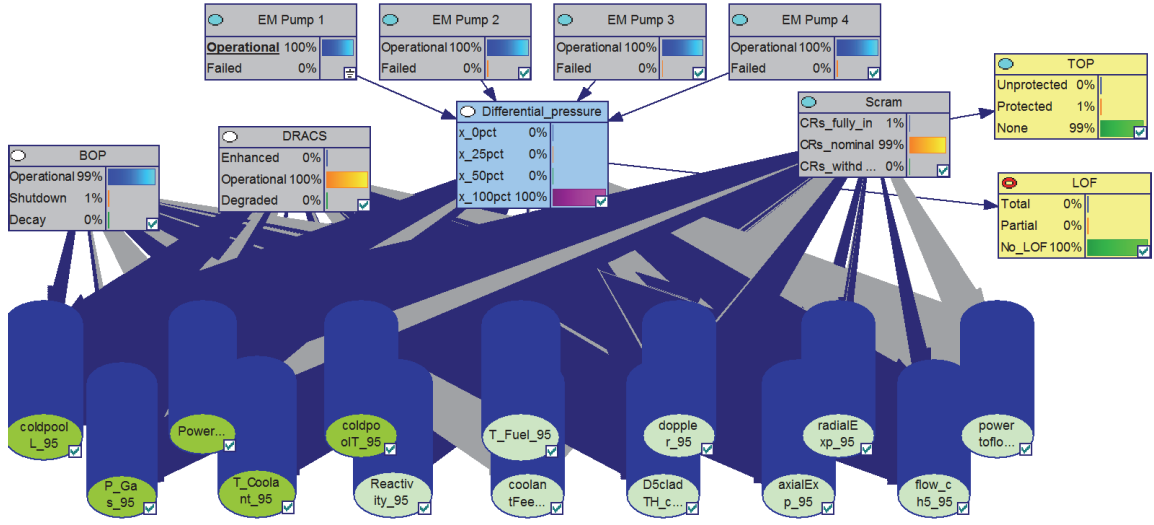

Figure 3: Unrolled DBN Model: the system creates a node for every time step. The model is now ready to be used to diagnose the conditions which lead to Transient Over Power (TOP) and Loss Of Flow (LOF) accident states.

information theory, the Kullback-Leibler divergence equation is used to measure the amount of information lost when the distribution function Q is used to approximate the actual distribution function P:

$$D(P||Q) = \sum_i P(i) log\left(\frac{P(i)}{Q(i)}\right) \quad (1)$$

where $P(i)$ represents the true probability distribution and $Q(i)$ represents a theoretical distribution. The equation defines divergence D between P and Q (i.e., the information lost when $Q(i)$ is substituted for $P(i)$).

In this study, K-L divergence is used to compare the BN model that includes all of the plant parameters with BN models that have removed one of the parameters. The divergence between two models shows how much information was lost by the elimination of one of the variables. If there is

a large amount of information lost when a node is removed, then the node is highly pertinent. If the information loss is minimal, than the node may be unnecessary and thus a candidate to be pruned.

In calculating the K-L divergence of an arc in our BN, $P(i)$ is the model with the node that we were measuring while $Q(i)$ was the model without the node that we were measuring. The values summed over i were combinations of possible observed and target states. K-L divergence is calculated for each arc between the observation and target nodes in a method similar to (Koiter 2006). Joint K-L divergence calculations are conducted over all the target nodes for each observation node. In calculating the joint K-L divergence, we treated each combination of possible target states as a single state in a joint target node that collected all targets into a single node.

Figure 4: The accuracies and F-scores of all target variable states with varied data sets. The baseline value is when the most frequent target combination is always chosen. Out of all the observed variables, $t\_coolant$ and $reactivity$ have the highest and lowest K-L divergence values respectively. Evaluating the BN while dropping each of these variables from the data set illustrates the change in performance of the BN's predictive power. For all of the target variables, dropping $reactivity$ has no effect. Whereas removing the coolant temperature reduces the BN's performance.

Table 1: The target variables and their states. Differential pressure is the difference in pressure between the electro-magnetic pumps. The $SCRAM$ parameter indicates the operators attempting to control the core's fission rate. Direct Reactors Auxiliary Cooling System (DRACS) is the plant's emergency heat removal system. Balance of Plant (BOP) consists of the systems that are not directly related to the nuclear steam supply systems.

Target	States	Prior Probabilities
Balance Of Plant	Operational	1.19×10^{-12}
	Shutdown	0.9999
	Decay	3.97×10^{-13}
Differential Pressure	0% flow	3e-13
	50% flow	1.04×10^{-4}
	100% flow	.9999
SCRAM	Fully In	0.0150
	Nominal	0.985
	Withdrawn	3.04×10^{-6}
DRACS	Enhanced	0.9850
	Operational	0.0150
	Degraded	7.95×10^{-12}

After calculating the K-L divergence of each of the observation variables we found that the coolant temperature ($t\_coolant$) had the highest K-L divergence value at 6.793×10^{-12} bits while $reactivity$ had the lowest K-L divergence value at 3.234×10^{-18} bits. Thus out of all of the observed variables, the $t\_coolant$ and $reactivity$ provide the highest and lowest amount of information gain respectively.

Cross-Validation

In any investigation of predictive models, cross-validation is a standard technique to measure the predictive power of a model while accounting for the possibility of an over-fit model (Ng 1997). Our training set, however, was taken from a spread of deterministically calculated simulations drawn from a *Discrete Dynamic Event Tree*. Unlike data drawn from a real source or random simulations, this data is representative of the real behavior of sodium reactor plants operating during beyond-design-basis accidents. It is our goal to measure how well the Bayesian network models the results of the SAS4a simulation.

To do this we measured the F-score and accuracy of each of the target variable states predicted by the Bayesian network when fed the discretized observation variable states produced by the SAS4a simulator. The F-score is useful for measuring the behavior of a machine learner on an asymmetric data set while the accuracy is generally used for symmetric data sets.

The results of the F-score and accuracy measurements can be found in Figure 4. The figure shows the F-score of the BN with full simulation data tends to outperform the baseline. In a few cases, however, this is not the case — Specifically, the F-score for DRACS $operational$ and the *SCRAM nominal* state. Both of these states are the chosen state for the baseline Bayesian network. This indicates that the Bayesian network is occasionally incorrectly predicting that these two states are not the correct state when they, in fact, are. This is to be expected when moving from a static baseline that always chooses the most likely target combination while ignoring the observation variables.

The accuracy of the DRACS operational and enhanced

states are the only states where the baseline guess is better than the BN with full observational data or observation data with dropped variables. There were also a few places where dropping the variable with the highest K-L divergence value actually improved results in the BN. These included two of the *SCRAM* states —- *Nominal* and *Withdrawn* in accuracy and the *SCRAM nominal* and BOP *Decay* on the F-score. This points to the importance of choosing the correct variables to predict the target accuracies on each of these variables. The highest K-L variable, coolant temperature ($t\_coolant$), was useful for predicting most variable states, but it caused a decrease in accuracy and F-score for some $SCRAM$ and BOP states. The differences in these scores shows how the Bayesian network compressed the SAS4a simulation results.

Conclusions and Future Work

This paper presented a prototype system for diagnosing the likely states of nuclear reactor systems. We simulated accidents with various conditions and used the resultant data to build a dynamic Bayesian network in order to learn the conditions of observed power plant parameters that can lead to *Transient Over Power* and *Loss Of Flow* Accidents. We also used the Kullback-Leibler divergence equation in order to gauge the relative importance of the plant's parameters.

This study examines the viability of the *Smart Procedures Dynamic Probabilistic Risk Assessments* methodology. This framework provides tools to diagnose the likely state of a nuclear reactor given the values of the observed plant parameters. This will potentially allow for a greater understanding of the state of a reactor during accidents when only a subset of information might be available. We also plan, based on Pearl's (2010) counterfactual analysis, to generate alternative paths forward, given a particular current state of the reactor. This understanding of possible alternative actions and their consequences will enhance operators' decision-making ability, especially during beyond-design-basis accidents.

Our current system uses the most basic Bayesian network techniques to capture the behavior of nuclear simulation results in a model that can be calculated in faster-than-real time. While it shows promising results, as it improves performance over our baseline analysis, we must continue to expand and calibrate the model parameters. In this paper we presented a model that used *Equal Width* discretization of observation variables. In the future we plan to use $Entropy$ or dynamic *Minimum Descriptive Length* partitioning (Clarke and Barton 2000) to discretize the observation variables. Fine tuning variables so that only those that have the best predictive power for each target is another important route forward for this research.

Acknowledgements

References

Cahalan, J.; Tentner, A.; and Morris, E. 1994. Advanced lmr safety analysis capabilities in the sassys-1 and sas4a computer codes. Technical report, Argonne National Lab., IL (United States). Funding organisation: USDOE, Washington, DC (United States).

Clarke, E. J., and Barton, B. A. 2000. Entropy and mdl discretization of continuous variables for bayesian belief networks. *International Journal of Intelligent Systems* 15(1):61–92.

Cowell, R. G. 2001. Conditions under which conditional independence and scoring methods lead to identical selection of Bayesian network models. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, 91–97. Morgan Kaufmann Publishers Inc.

Druzdzel, M. J. 1999. SMILE: Structural modeling, inference, and learning engine and GeNIe: a development environment for graphical decision-theoretic models. In *Proceedings of American Association for Artificial Intelligence (AAAI-99)*, 902–903.

FNI. 2012. Investigation committee on the accident at fukushima nuclear power stations. Technical report, Final Report, July 23, 2012.

Groth, K. M.; Denman, M. R.; Cardoni, J. N.; and Wheeler, T. A. 2014. 'Smart Procedures': Using dynamic PRA to develop dynamic, context-specific severe accident management guidelines (SAMGs). In *Proceedings of the International Conference on Probabilistic Safety Assessment and Management (PSAM 12)*.

Groth, K. M.; Denman, M. R.; Jones, T. J.; Darling, M. C.; and Luger, G. F. 2015. Proof-of-concept accident diagnostic support for sodium fast reactors. In *European Safety and Reliability Conference ESREL 2015, Zurich Switzerland*.

Kim, K., and Bartlett, E. B. 1996. Nuclear power plant fault diagnosis using neural networks with error estimation by series association. *Nuclear Science, IEEE Transactions on* 43(4):2373–2388.

Koiter, J. R. 2006. Visualizing inference in Bayesian networks. Master's thesis, Delft University of Technology.

Murphy, K. P. 2002. *Dynamic Bayesian networks: representation, inference and learning*. Ph.D. Dissertation, University of California, Berkeley.

Ng, A. Y. 1997. Preventing "overfitting" of cross-validation data. In *ICML*, volume 97, 245–253.

Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann.

Pearl, J. 2010. *Causality*. Cambridge university press.

Yang, Y., and Webb, G. I. 2002. A comparative study of discretization methods for naive-bayes classifiers. In *Proceedings of PKAW*, volume 2002. Citeseer.