HEALTH MONITORING FOR COMMERCIAL AIRCRAFT SYSTEMS

K. Wojtek Przytula*, David Allen*, John Vian, Gary Mansouri****
***HRL Laboratories, LLC, Malibu, CA 90265**
****The Boeing Company, Phantom Works, Seattle, WA 98124**

Keywords: *aircraft health monitoring, data mining, probabilistic models*

Abstract

This paper presents a methodology for the development of health monitoring solutions for complex systems. The goal of monitoring is early detection of system failures based on present and historic values of signals from built-in sensors and detectors. The methodology is illustrated on an example of health monitoring for an aircraft bleed air control system. The monitoring solution was developed using more than 100 signals collected from a real commercial aircraft over more than 50 flights. The system failures are detected using reasoning on a graphical probabilistic model; more specifically a Bayesian network. The model was created from both knowledge of the bleed air control system and learned from the data. The operation of the monitoring software was verified on test data and resulted in more than a 98% success rate of failure identification.

1 Introduction

Advanced complex systems, such as commercial aircraft systems, consist of a very large number of components which closely interact with each other. As the cost of electronic and computer hardware decreases, the systems are equipped with increasing numbers of sensors, detectors and computerized controllers. The additional hardware makes it possible to monitor the health of the systems with improving accuracy. However, the complexity of health monitoring solutions also rapidly increases. It therefore becomes necessary to use advanced methodology and software tools to develop effective monitoring systems.

System health monitoring is a form of system diagnosis, in which the goal is to detect system failure and identify, which component is responsible for it. In monitoring, the diagnosis is based only on observations derived from signals originating from built-in sensors and detectors, e.g. pressure sensors or valve position detectors. It does not take into account the symptoms of failure, e.g. abnormal sounds or vibrations, or measurements performed by means of external devices, such as portable testers, which are often used during troubleshooting of the aircraft systems in repair shops. Although health monitoring is limited to the built-in devices, it has an advantage of providing real-time health status either during the flight and/or soon after its completion. It is very useful for a go-no-go decision at the airport gate and may be of critical importance in decisions affecting loss of human life and/or damage to expensive hardware.

We illustrate the application of system health monitoring on an example of a commercial aircraft air bleed control system [1]. The system provides air to several other aircraft systems including the passenger cabin air conditioning system. There are over a hundred different signals available in a typical commercial aircraft, which are of potential utility in monitoring this system's health. Real-time monitoring of the signals results in tens of thousands of data records per flight.

This paper describes a methodology and supporting software tools used for the development of monitoring solutions for complex systems. Our methodology involves two steps: derivation of diagnostic observations and creation of diagnostic models. The observations are derived from the sensor and detector signals. We begin with identifying these signals that individually or in combination provide an indication of the component failure. Next, we develop algorithms, which take in the selected signals and produce diagnostic observations.

The second step is the creation of a diagnostic model. We use a form of graphical probabilistic models – Bayesian networks [2]. The models capture relations between diagnostic observations and component failure modes. We use the models and a probabilistic reasoning engine to derive the likelihood of component failure given the state of the diagnostic observations.

For efficient development of diagnostic observations and creation of the diagnostic models we use a number of commercial and custom software tools. The tools are necessary to be able to handle the large numbers of signals and data records.

Some examples of other work that have focused on model based diagnosis of aircraft systems include [3]-[8]. A more thorough analysis of the Bayesian approach to diagnosis and prognosis has also been undertaken [9]. Model based approaches and other related pattern recognition methods have also been applied to many other domains [10]-[13].

2 Methodology for Development of Health Monitoring Systems

This section is devoted to the discussion of our methodology for development of health monitoring solutions for complex systems. We will also describe the software tools which support our methodology.

We assume that we are developing health monitoring for an existing system, for which the sensors and detectors have already been selected. Also the sampling and collection of signal data has been determined.

2.1 Procedure for Monitoring System Development

The development of health monitoring solutions begins with the collection of data and knowledge about the system. The data are sampled values of all the pertinent signals over an extended period of time. For an aircraft system the data for tens to hundreds of flights are needed. We also expect that the data contain signals documenting failure modes of the system components and that there are annotations indicating when and what failure occurred. The data on component reliability are also very beneficial. The information about the system would typically include a diagram or schematics and a functional description. Alternatively system knowledge may be acquired directly from an expert.

Our approach to health monitoring requires definitions of the diagnostic observations and the diagnostic model, Fig. 1. They are obtained from data and domain knowledge. The diagnostic observations are computed using one or more signals from the sensors and detectors. The computations extract from the raw signals the information useful for diagnosing component failures. A simple example of such a processing is smoothing of a signal by filtering, followed by comparison of the value to a predefined threshold. The observation derived from the signal may take two states: “high” – when the filtered signal is above the threshold and “normal” – when it is below the threshold.

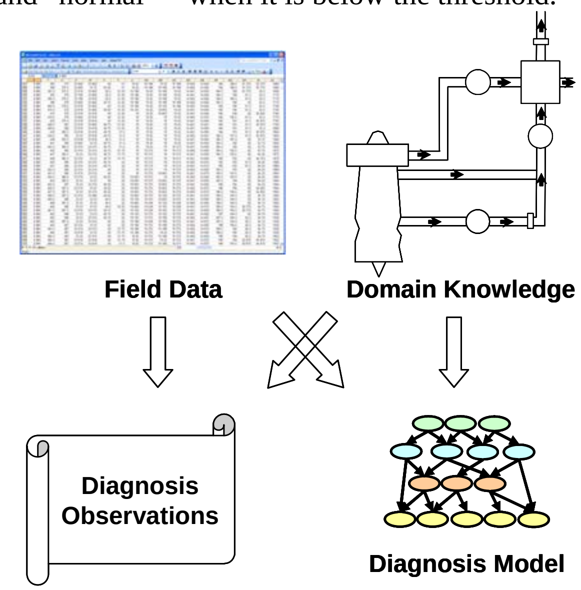

Fig. 1. Development of a health monitoring solution for a complex system.

We use Bayesian networks for our diagnostic model, [2]. The model is an annotated graph, whose nodes represent elements of the domain, i.e. components, systems, diagnostic observations, and measures of usage, Fig. 2. The directed links between the nodes encode relations, i.e. a link between given component node (a parent) and observation node (a child) indicates that failures of the component result in a change of the state of observation. The annotations are conditional probabilities, which represent the strength of the relations. We use layered Bayesian networks for diagnosis, because they are easier to create and are less demanding computationally [14].

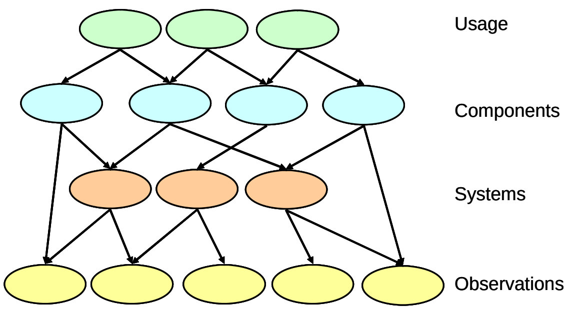

Fig. 2. Layered Bayesian network model for diagnosis.

The diagnostic model is then used to obtain the probability of component failure given the states of the diagnostic observations. The model represents a joint probability distribution Pr over the variables X_1 , X_2 , …, X_n , which according to the chain rule is computed as:

$$\Pr(X_1, X_2, \ldots, X_n) = \Pr(X_n | X_{n-1}, \ldots X_2, X_1) * \ldots * \Pr(X_2 | X_1) * \Pr(X_1) \qquad (1)$$

For a Bayesian network this rule can be written as:

$$\Pr(X_1, X_2, \ldots, X_n) = \prod_{i=1}^{n} \Pr(X_i | \mathbf{Pa}_i) \qquad (2)$$

where $\mathbf{Pa}_i$ represents all parent nodes of the node X_i .

Once the algorithms for diagnostic observations and the diagnostic model are developed, we are ready to implement the complete health monitoring system, Fig. 3. The system takes in the sensor and detector signals and produces from them the diagnostic observations. The observation states and the diagnostic model are combined using a probabilistic reasoning engine, which implements formulas such as (2). The engine produces the probability of component failure given the observation states, i.e. system diagnosis.

We will consider two possible scenarios for health monitoring. One is a real-time monitoring, in which a new sample of signals is processed as soon as it is available and updated health results are immediately available. The other is batch processing, in which data are collected over an extended period of time and the results are computed for all the collected data. In the case of aircraft health monitoring, the batch results could be available at the end of a flight phase, e.g. take off, or at the end of the entire flight. The choice of the scenario depends on the monitoring requirements for the specific system, as well as capabilities of the on-board hardware.

General similarities can be drawn between our proposed methodology for health monitoring systems and the process of knowledge discovery in databases [15]. For example, one can see how our derivation of diagnostic observations ties in with the KDD process of data reduction/projection, with the end goal of both methodologies to build a model capable of efficiently analyzing/monitoring data.

2.2 Software Tools for Health Monitoring System Development

Efficient development of the diagnostic observations and the diagnostic models for a complex system requires software tools. The tools help in manipulating the large data sets, in developing observation algorithms and diagnostic models as well as evaluating the different algorithms and models.

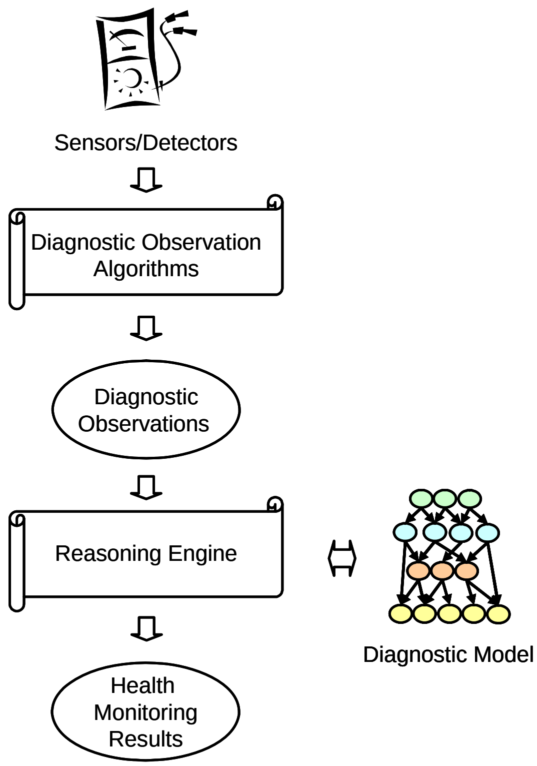

Fig. 3. Implementation of a health monitoring solution for a complex system.

Our first challenge is to select signals that are pertinent to health monitoring of our system from all the signals available to us. In the selection we need to use both the understanding of the system and of the data. The understanding of system operation helps in focusing on a candidate subset of signals. The subset may include signals that appear unrelated, but may be useful in detecting abnormal system behavior, e.g. equivalent signals for another aircraft engine.

Our understanding of the data can be significantly improved by visualization of the signals with the failure annotations. The visual inspection also helps in identifying errors and noise in the data, e.g. dropped signals, spikes, etc.

The visualization can be implemented in a commercial tool such as Mathworks Matlab. For manipulation of the data, i.e. selection of individual signals and fragments of signal history, we will need a database and database management tools such as Microsoft SQL Server.

The cleaning of data and preprocessing for visualization may be implemented using the database tools as well as off-the-shelf and custom data mining tools, e.g. Microsoft's Data Mining Tools. These tools contain routines such as min, max, average and various forms of filtering. To develop diagnostic observation algorithms we need to be able to process and visualize multiple signals at a time. Here custom algorithms may be most suitable.

Development of diagnostic models requires still different tools. Some commercial tools provide support for generation of simple models from data; for example decision trees or naïve Bayesian networks can be created using Microsoft's Data Mining Tools. However, to combine domain knowledge with the data more sophisticated tools are required. HRL has developed a family of tools for creation, debugging, evaluation, and updating of layered Bayesian networks from data [14],[16]-[17].

Our work on the development of system monitoring solutions involved integration of several commercial and custom software tools as depicted in Fig. 4.

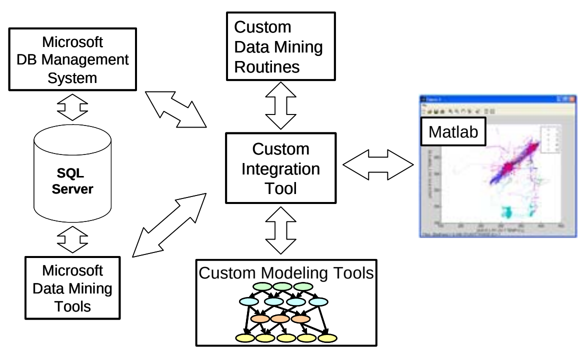

Fig. 4. Software tools for development of health monitoring solution.

3 Application Example

3.1 Aircraft Bleed Air Control System

In this section we apply our methodology to build a health monitoring solution for a real-world domain, that of the bleed air control system (BACS) of a commercial aircraft. This system is part of the environmental control system (ECS) and is responsible for providing air at the appropriate pressure and temperature for use by other aircraft systems, including the passenger cabin.

The first step in the process is to understand the problem domain. Fig. 5 depicts a simplified schematic of the bleed air control system of the commercial aircraft we studied. More detailed information on the BACS system in another related aircraft can be obtained from [1]. In the diagram, hot air is extracted from the aircraft's engine through either the intermediate stage or high stage, depending on the setting of the high pressure valve. This air then passes through another safety valve to a mechanical heat exchanger called the precooler. In addition to the hot air, cold air is also extracted from the engine and regulated by another valve before reaching the precooler. This cold air is used to regulate the temperature of the hot air resulting in an air supply for other aircraft systems at a safe temperature and pressure. There are many feedback, control, and safety mechanisms in place in this system. Therefore, domain knowledge is especially important for accurate diagnosis, as simple analysis of the observations may not be sufficient, if the redundancy and control mechanisms are initially able to compensate for a failure.

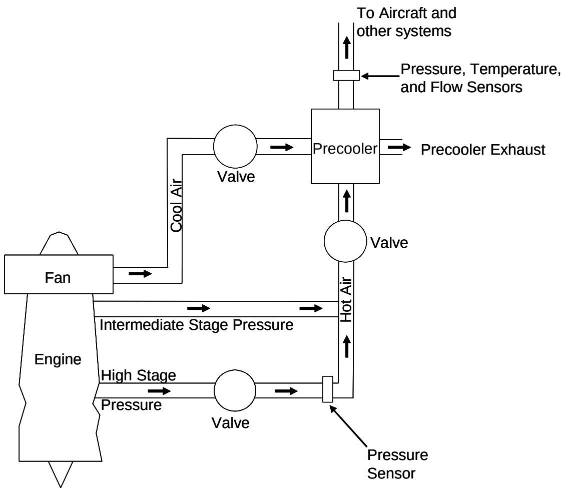

Fig. 5. Diagram of a commercial aircraft bleed air control system.

Some previous work has also been done examining this aircraft system [18]-[21]. This previous work has examined methods for optimal troubleshooting and analysis of system reliability and maintainability.

The second step is to create or select a target dataset. The data used in our experiments were collected from a real commercial aircraft and consist of a subset of raw signals (we examined more than 100 signals) and over 50 flights worth of data (hundreds of thousands of records for each signal). These signals were annotated with timestamps, allowing flights to be distinguished, and additionally includes the time of the replacement of one of the system components.

In our work, we examined this dataset and used it to derive diagnostic observations. We then used some of the data from these observations to train our diagnostic model, while the remainder was used for model validation. The final output of the project was a solution for system health monitoring.

3.2 Derivation of Diagnostic Observations

The first step in our derivation of diagnostic observations, after importing our data into a database management system, was to use domain knowledge and simple signal statistics to select an initial subset of signals which may be of direct relevance to the system. These included many signals shown in Fig. 5, such as the temperature and pressure signals and the states of the valves. In addition, other signals from the engine and flight data such as speed and altitude were included. We also used system redundancy and symmetry. For example we compared signals from two different engines and analyzed differences in pressures from different points within the BACS system.

During this stage we also analyzed the signals for unusual changes across individual flights or across multiple flights. These were sometimes determined to be valid diagnostic observations and sometimes were indicative of noise in the signal. We therefore leveraged the ability of the database system and created a general filtering mechanism to remove data during signal loss or out of range values.

We then visualized individual signals in order to analyze their changes and see their correspondence with the external annotations. In addition to analyzing and visualizing entire flights we also partitioned flights into segments and analyzed different flight segments such as takeoff, cruising, and landing. We then compared data attributes such as minima, maxima, and averages across flights and flight segments.

In addition to analyzing and visualizing individual signals and signal statistics we also developed methods for analyzing pairs of signals, as shown in Fig. 6. In this figure we depict the interdependence of two signals in a 2D XY plot (the different types of data points depicted in the legend will be explained in the next section). We also incorporated methods for visualizing subsets of signals which depend very heavily on one another (e.g. the pressure signals are very closely tied to the settings of the valves).

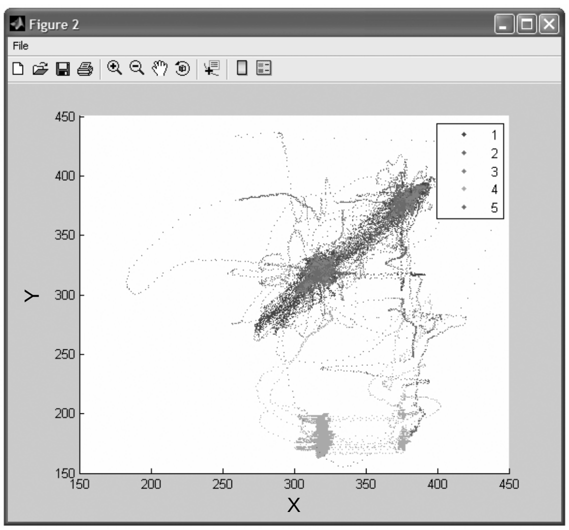

Fig. 6. XY Plot of two related signals.

3.3 Learning of the Diagnostic Model

Through the data analysis and visualization techniques discussed above we were able to define distinct classes (or modes of operation) which we were interested in distinguishing. Differentiating these classes then became the goal of the health monitoring solution. In Fig. 6, the legend depicts 5 modes of operation and each mode is plotted using a distinct color allowing the user to observe patterns in the diagnostic observations' interaction related to the mode of operation. In preparation for building and validating our classifier model we split the data into a set of training data (75%) and a set of testing data (25%).

The next steps of the process are to select specific algorithms and build the models. We used the training data to learn the parameters of many different classification models. For example, we examined different types of classifiers such as Naïve Bayes networks and decision trees. Additionally we also examined different combinations of possible diagnostic observations. In this paper we further discuss the Naïve Bayes models we developed.

Naïve Bayesian classifiers [22] make strong assumptions on the independence of observations, however they have been used very successfully in practice [23]. In Fig. 7 we depict a simplified version of our model. In the center of the figure you can see a variable depicting the system health (the mode of system operation). Surrounding that in the figure is a subset of the relevant diagnostic observations used in our model. The model also consists of a set of parameters for each edge specifying the weight of the dependence.

Once we learned the model from our training data we were able to analyze the sensitivity of the model to individual diagnostic observations. These techniques allow for distinguishing observations which are most relevant and determining how sensitive the model is to changes in parameters and observations [24]-[25]. This allowed us to further reduce our set of diagnostic observations and iteratively refine our model and analyze its robustness.

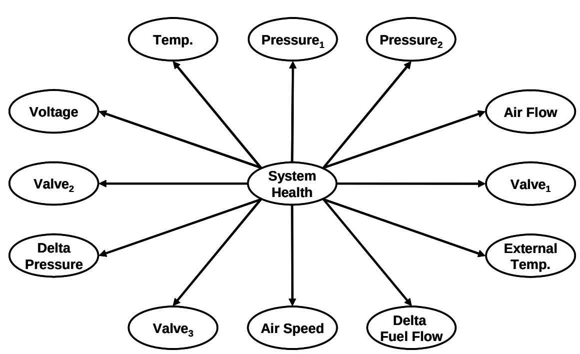

Fig. 7. Simplified Bayesian network diagnosis model obtained by learning from real aircraft data.

3.4 Failure Detection Results

Once our health monitoring solution was created we used it to examine the testing data to validate the model and interpret the results. Our testing data consisted of 25% of the initial data set, which was still quite large (hundreds of thousands of records). We then classified each record into the different modes of operation and compared that with our labeled ground truth.

The resulting confusion matrix is shown in Table I. In the table, each column corresponds to one of the classes in the ground truth, while each row corresponds to the values predicted by the model. The classes include one for normal operation and three for abnormal operation. As you can see, the model was very accurate at predicting the correct class (values along the diagonal). For example, it was able to accurately classify all samples which were actually class 1. It was most inaccurate (98.1%) predicting class 3, because the diagnostic observations are very similar to class 1. Overall the model was 99.8% accurate at distinguishing between these classes on the testing data.

The final step we undertook was to formalize an interface to the system health monitoring solution and to demonstrate the technology to potential end users.

Table I.
Diagnosis results obtained using the diagnostic model and real test data for aircraft system.

	1 (Actual)	2 (Actual)	3 (Actual)	4 (Actual)
1 (Predicted)	100.0%	0.2%	1.7%	1.2%
2 (Predicted)	0%	99.6%	0.2%	0.1%
3 (Predicted)	0%	0.2%	98.1%	0%
4 (Predicted)	0%	0%	0%	98.7%

4 Conclusions and Future Work

Development of health monitoring solutions for complex systems requires a systematic approach and use of software tools. The tools are indispensable for manipulation of large sets of field data and assist in derivation of diagnostic observations from raw sensor and detector signals. They are also necessary for efficient development of accurate diagnostic models.

The methodology and tools described in this paper offer an example of an approach that has been verified on the development of a real monitoring solution for a complex aircraft system and the use of a large set of real field data. We believe that the methodology worked well and produced an accurate and efficient solution. The toolset put together during the project can be easily reused for other similar systems, offering an environment for rapid development of new solutions.

We are planning to continue the work described in this paper. In particular we would like to extend our results from system health diagnosis to health prognosis. System health prognosis is a much harder task and will require more advanced techniques, e.g. models based on dynamic Bayesian networks, etc., and significantly more field data.

Acknowledgement

The authors would like to acknowledge Tsai-Ching Lu and Arthur Choi for their technical suggestions and Danny Dever for providing access to data and feedback on the results.

References

[1] E. Hunt, D. Reid, D. Space and F. Tilton, “Commercial Airliner Environmental Control System,” Aerospace Medical Association annual meeting, 1995.

[2] J. Pearl, “Probabilistic Reasoning in Intelligent Systems,” Morgan Kaufmann Publishers, Inc., San Francisco, CA, 1988.

[3] C. Byington, P. Kalgren, R. Johns and R. Beers, “Embedded Diagnostic/Prognostic Reasoning and Information Continuity for Improved Avionics Maintenance,” Autotestcon, 2003.

[4] C. Byington, P. Kalgren, B. Dunkin, and B. Donovan, “Advanced Diagnostic/Prognostic Reasoning and Evidence Transformation Techniques for Improved Avionics Maintenance,” Proceedings of the 2004 IEEE Aerospace Conference, Vol. 5, 2004.

[5] S. Letourneau, S. Famili and S. Matwin, “Discovering Useful Knowledge from Aircraft Operation/Maintenance Data,” Proceedings of the Workshop on Machine Learning in the Real World, 1997.

[6] S. Letourneau, F. Famili and S. Matwin, “Data Mining for Prediction of Aircraft Component

Replacement," IEEE Intelligent Systems Journal, 1999.

[7] J. Pena, F. Famili, and S. Letourneau, "Data Mining to Detect Abnormal Behavior in Aerospace Data," International Conference on Knowledge Discovery and Data Mining, 2000.

[8] K. Keller, K. Swearingen, J. Sheahan, M. Bailey, J. Dunsdon, K.W. Przytula, and B. Jordan, "Aircraft Electrical Power Systems Prognostics and Health Management," Proceedings of the 2006 IEEE Aerospace Conference, 2006.

[9] J. Sheppard and M. Kaufman, "A Bayesian Approach to Diagnosis and Prognosis Using Built-In Test," IEEE Transactions on Instrumentation and Measurement, Vol. 54, No.3, 2005.

[10] S. Jin, D. Yeung, X. Wang, and E. Tsang, "A Feature Space Analysis for Anomaly Detection," IEEE International Conference on Systems, Man, and Cybernetics, 2005.

[11] R. Hewett, "Assessment of Software Risks with Model-Based Reasoning," IEEE International Conference on Systems, Man, and Cybernetics, 2005.

[12] M. Muguira, J. Salton, D. Novick, J. Schwebach, and C. Wilson, "Chile Identification for Metrics in the Chile Industry," IEEE International Conference on Systems, Man, and Cybernetics, 2005.

[13] W. Melek, A. Sadeghian, H. Najjaran, and M. Hoorfar, "A NeuroFuzzy-based Expert System for Disease Diagnosis," IEEE International Conference on Systems, Man, and Cybernetics, 2005.

[14] T.C. Lu and K.W. Przytula, "Methodology and Tools for Rapid Development of Large Bayesian Networks," 16th International Workshop on the Principles of Diagnosis (DX-05), 2005.

[15] U. Fayyad, G. Piatetsky-Shapiro, and P. Smyth, "From Data Mining to Knowledge Discovery in Databases," AI Magazine, Fall, 1996.

[16] K.W. Przytula, D. Dash, and D. Thompson, "Evaluation of Bayesian Networks under Diagnostics," Proceedings of the 2003 IEEE Aerospace Conference, 2003.

[17] K.W. Przytula, G. Isdale, and T.C. Lu, "Collaborative Development of Large Bayesian Networks," Proceedings of the 2006 AUTOTESTCON, 2006.

[18] G. Clark and R. Paasch, "Diagnostic Modeling and Diagnosability Evaluation of Mechanical Systems," ASME Journal of Mechanical Design, Vol. 118, No. 3, 1996.

[19] M. Murphy and R. Paasch, "Reliability Centered Prediction Technique for Diagnostic Modeling and Improvement," Research in Engineering Design, 1997.

[20] M. Fitzpatrick and R. Paasch, "Analytical Method for the Prediction of Reliability and Maintainability Based Life-Cycle Labor Costs," ASME Journal of Mechanical Design, Vol. 121, No. 4, 1999.

[21] R. Paasch and P. Durgi, "Optimal Troubleshooting for Electro-Mechanical Systems," Proceedings of the 2003 ASME Design Automation Conference, 2003.

[22] I. Rish, "An empirical study of the naïve Bayes classifier," International Joint Conference on Artificial Intelligence, Workshop on Empirical Methods in Artificial Intelligence, 2001.

[23] P. Domingos and M. Pazzani, "On the optimality of the simple Bayesian classifier under zero-one loss," Machine Learning, 1997.

[24] H. Chan and A. Darwiche, "Sensitivity Analysis in Bayesian Networks: From Single to Multiple Parameters," Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence, 2004.

[25] O. Kipersztok and G. Dildy, "Evidence-based Bayesian Networks Approach to Airplane Maintenance," International Joint Conference on Neural Networks, Vol. 3, 2002.

Copyright Statement