RESULTS OF A PROBABILISTIC FAULT DETECTION AND DIAGNOSIS METHOD FOR VAPOR COMPRESSION CYCLE EQUIPMENT

AUTHORS

Margaret B. Bailey, Ph.D., P.E., Associate Member ASHRAE
Jan F. Kreider, Ph.D., P.E., Member ASHRAE
Peter S. Curtiss, Ph.D., Member ASHRAE

KEYWORDS

air conditioning, chiller, classification, compressor, computer, laboratory, leakage, oil, refrigerant, research

ABSTRACT

Reliable, automated fault detection and diagnosis of vapor compression cycle equipment is a valuable feature for equipment owners due to energy and environmental concerns. This paper describes designing and testing the viability of a fault detection and diagnosis system for chillers. The fault detection and diagnostic (FDD) tool performs the task of analyzing archived chiller operating data and detecting faults through recognizing patterns within the data. The tool utilizes a neural network program trained on chiller operating data collected during normal and fault conditions at a full-scale heating, ventilation and air conditioning laboratory. Faults include the improper charging of refrigerant and oil as well as air-cooled condenser fouling and fan loss. The final, trained FDD tool is capable of correctly predicting the current state of chiller operation with a misclassification rate of three percent.

AUTHORS NOTE

M.B. Bailey is Assistant Professor, Department of Civil and Mechanical Engineering, United States Military Academy, West Point, New York; J.F. Kreider is Professor and Founding Director, Joint Center for Energy Management, Department of Civil, Environmental and Architectural Engineering, University of Colorado, Boulder and President of Kreider & Associates, LLC., Peter S Curtiss is an Adjunct Professor, Joint Center for Energy Management, Department of Civil, Environmental and Architectural Engineering, University of Colorado, Boulder and a Consulting Engineer at Kreider & Associates, LLC .

RESULTS OF A PROBABILISTIC FAULT DETECTION AND DIAGNOSIS METHOD FOR VAPOR COMPRESSION CYCLE EQUIPMENT

ABSTRACT

Reliable, automated fault detection and diagnosis of vapor compression cycle equipment is a valuable feature for equipment owners due to energy and environmental concerns. This paper describes designing and testing the viability of a fault detection and diagnosis system for chillers. The fault detection and diagnostic (FDD) tool performs the task of analyzing archived chiller operating data and detecting faults through recognizing patterns within the data. The tool utilizes a neural network program trained on chiller operating data collected during normal and fault conditions at a full-scale heating, ventilation and air conditioning laboratory. Faults include the improper charging of refrigerant and oil as well as air-cooled condenser fouling and fan loss. The final, trained FDD tool is capable of correctly predicting the current state of chiller operation with a misclassification rate of three percent.

INTRODUCTION

The FDD method developed here incorporates a neural network (NN) classifier to infer the current state of a chiller given a vector of observables that consists of various operating properties. The method relies on the availability of normal and fault data for training purposes. Therefore, a fault library was assembled for this purpose as discussed in Bailey et al. (2000). The experimental data was gathered at a full-scale heating, ventilation, and air conditioning (HVAC) laboratory (Kreider et al. 1999). The facility encompasses a total area of approximately 7,500-square-feet that includes two, 400 square-foot full size zones. The remainder of the laboratory houses air handling units, zone simulators, a boiler and a chiller plant. All of the equipment is full-scale and the laboratory's cooling load originates from conditioning large quantities of preheated outdoor and/or return air. The magnitude and profile of the laboratory's cooling load are user-defined and the desired outdoor air condition can be simulated.

The chiller plant consists of a commercially available 70-ton helical rotary-screw, air cooled chiller with a remote 50-ton air-cooled refrigerant condenser. The chiller plant provides chilled water for the cooling coils within the lab. It consists of one direct expansion evaporator with a dual circuit configuration and two independent circuits each including a screw compressor, oil cooler, filter drier, sight glass, electronic expansion valve, and charging valve. The air-cooled condenser includes six constant speed, vertical discharge, direct-drive fans.

In order to collect data internal to the chiller, the first author installed a new data acquisition system. The data collection setup includes a manufacturer supplied hardware board and software program. Through the software, a list of tokens or points is user defined. The hardware board and software combination detects the specified tokens on the chiller communication bus via a twisted communication pair. Data are recorded into predefined data files according to user defined triggers and at a user specified time interval of six seconds.

NN applications include both regression and classification or pattern recognition (Massie 2000). The NN classifier described in this paper is used for pattern recognition. It detects complex relationships and trends that exist within the input space and uses the relationships to classify the current state of the system. Several complex trends and patterns exist within the operating data as discussed in Bailey (1998). The FDD system design objectives include low false alarm rate, passive performance, and the ability to correctly classify faults in both steady state and transient chiller operation.

FDD can occur in either real-time or near real time. Although both methods have their associated advantages, near real time FDD is adopted in this research. Near real time FDD is performed off-line at a remote personal computer or workstation with limited memory requirements. This approach is quite effective for off site service contracts as well as future revisions, improvements, and maintenance of the FDD program. In addition, most modern chillers possess the capability of archiving operating data off-line using the on-board control packages. Using this option, operating data are sent to a remote computer. The data are archived and evaluated on a regular, predefined basis by the FDD program. Removal of past data is performed regularly by the program.

Detection of the faults considered in this research is possible by a human expert, however individuals with this expertise are often unavailable or unable to closely monitor a given piece of equipment. Therefore, a computer program, which adopts a consistent approach, is developed to successfully perform the difficult task of fault detection or pattern recognition. A natural and well-established setting to find an answer to pattern recognition problems is by using a statistical approach. The basic premise to pattern recognition is that the operating data exhibits an underlying systematic aspect that may be corrupted with random noise. The goal of the classifier is to map the underlying systematic aspects of the data without capturing the noise component (Bishop 1995). With this goal achieved, the classifier possesses the ability to generalize or correctly classify previously unseen data or features. A classic application of pattern recognition is FDD. In the following paper, several different methods of pattern recognition are employed to detect and diagnose faults within chillers.

CHILLER FAULTS CONSIDERED

An informal survey was conducted early in the research process to identify the most important faults associated with rotary screw chillers. The information collected includes fault type, frequency of occurrence and the total annual costs associated with the repair of each fault. Seven possible and probable states of chiller operation were determined after numerous discussions with representatives from a major chiller manufacturer and a large chiller service contractor. The chiller manufacturer representatives involved in these discussions included engineers and managers involved in the development of chiller compressors and controls. The chiller faults that were deemed important and significantly common include refrigerant over and under charge, oil over and under charge, condenser fouling and loss of an air-cooled condenser fan.

PRELIMINARY CLASSIFIER DESIGN

Two basic classification methods are first studied and evaluated to determine the factors that significantly affected classifier performance. The basic classifier designs include the non-parametric K-nearest neighbors (KNN) classifier and a linear classifier minimized using the mean squared error (MSE) Criterion. For more information regarding these classification methods consult Bishop (1995).

The database used by the simple KNN and MSE classifiers consists of chiller plant data gathered over a two-month period. The initial stages of data preprocessing involves reducing the high dimensionality of the input space by discarding a subset of the original inputs. Through expert knowledge, the input space's dimensionality is reduced to seven features without the loss of significant information. The final seven input features include:

- average phase current
- outdoor air temperature
- evaporator entering water temperature
- compressor refrigerant suction temperature
- saturated refrigerant evaporating temperature
- refrigerant condensing temperature
- entering compressor oil temperature

The chiller plant was operated in seven modes while data are collected. Figure 1 is a schematic diagram of the classifiers' inputs and outputs.

INPUT SPACE

1. outdoor air temp
2. evaporator ewt
3. suction temp
4. sat evap temp
5. cond temp
6. avg motor phase current
7. enter comp oil temp
8. time

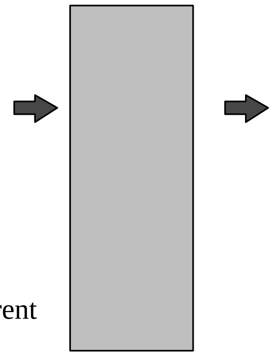

POSTERIOR PROBABILITIES

1. normal
2. refrigerant undercharge
3. refrigerant overcharge
4. oil undercharge
5. oil overcharge
6. condenser fan loss
7. condenser fouling

FIGURE 1. Typical Schematic of Classifiers with Input and Output Information

In addition to dimensionality reduction, data preprocessing also involves normalizing all features and clustering or binning the data according to time. Therefore, all the chiller data are categorized according to time of day. Most of the data files are recorded over the same time interval each day, namely 7:00 AM to 3:00 PM. Therefore, the data sets are divided into bins according to the minutes into each day. Two forms of binning are tested, namely linear and logarithmic. Linear binning includes separating the data sets into six equal time groups, approximately 60 minutes in length each. Logarithmic binning involves separating the data into ten non-equal time groups. The width of the first time bin was considerably smaller than that associated with the last time bin. Each successive bin width was increased by a factor of 1.33. The length of each logarithmic bin ranged from 7.3 minutes for the first bin to 75.5 minutes for the last bin. This binning scheme was attempted to account for the transient start-up conditions displayed in the beginning of the data set.

A summary of the misclassification rate associated with each classifier and each binning method is listed in Table 1. The misclassification rate listed in Table 1 was determined using the following relationship.

$$Misclassification\ Rate = \frac{(number\ of\ misclassified\ samples)}{total\ number\ of\ testing\ patterns} \qquad (1)$$

Note that the overall performance of the two classifiers is similar, however during certain bins, one classifier outperforms the other. For example, for linear bin six, the KNN classifier dramatically outperforms the MSE classifier due to the nature of the data included within this time bin. From the results listed in Table 1, it appears that the binning technique most greatly affects the classifier's performance however neither classifier performed acceptably. The classifiers' performance problems in time bin one can be attributed to the start-up dynamics of the chiller. Neither classifier was capable of adapting to these dynamics. The source of bin eight's classification problems is most likely associated with the cycling of the chiller stages at this point of experimental testing and the width of this time bin.

TABLE 1.

Misclassification Rates Associated with KNN and MSE Classifiers Using Linear Binning (60 Minute Intervals) and Logarithmic Binning (7.3 to 75.5 Minutes Intervals)

Linear		**Bin Numbers**									
Binning		**1**	**2**	**3**	**4**	**5**	**6**				**AVG.**
	KNN	52%	43%	48%	34%	26%	15%				**36%**
	MSE	49%	27%	37%	34%	27%	30%				**34%**
Log		**Bin Numbers**									
Binning		**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**AVG.**
	KNN	43%	9%	14%	24%	20%	24%	35%	41%	22%	13%	**25%**
	MSE	51%	16%	16%	29%	27%	19%	30%	35%	26%	23%	**27%**

The results of this exercise highlight several important factors that are used in the design of a non-linear, NN classifier. First, the performances of linear and non-parametric classifiers are demonstrated to be inadequate for the application of chiller FDD. The rates indicated in Table 1 are exceedingly high and unacceptable. Furthermore, the issue of time binning appeared to be critical. Finally, the transient behavior during the chiller start-up period results in poor performance by each classifier regardless of binning approach.

NEURAL NETWORK CLASSIFIER

The results from the preliminary classifiers provide insight into the problem of FDD of chillers that is very useful in the design of a NN classifier. NNs are well suited for pattern recognition due to their ability to establish non-linear decision boundaries between classes as shown in Figure 2. The advantages of NN models compared with other modeling techniques include their ability to provide a very general and powerful framework for representing non-linear mapping. Barron (1993) showed that NNs offer a dramatic advantage for function approximation in spaces of many dimensions. In addition, if designed correctly, NN classifiers can predict probability of class membership.

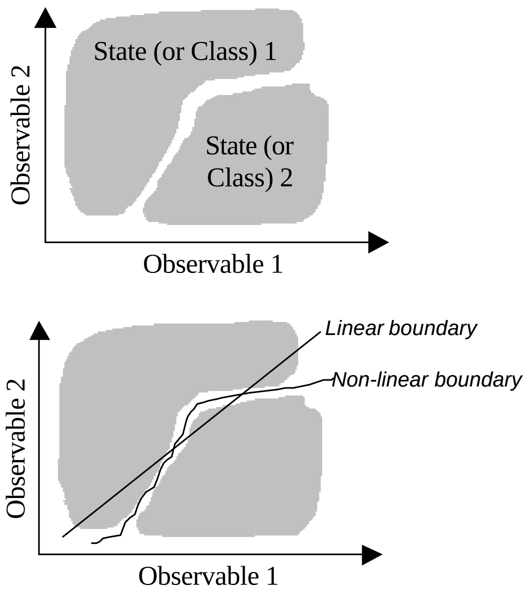

FIGURE 2. Advantage of Using a Non-Linear Decision Boundary for Classifying Two Dimensional Data

The first stage in any classification process is inference. In this research, the NN model uses an input vector, that consists of twenty-eight variables, to determine or infer values for the posterior probabilities, $P(C_i|\underline{x})$ as shown in Figure 3. The posterior probability represents the probability that class C_i exists given the observable, $\underline{x}$. The posterior probability is then used in the second stage of classification called the decision making process as discussed later in this section.

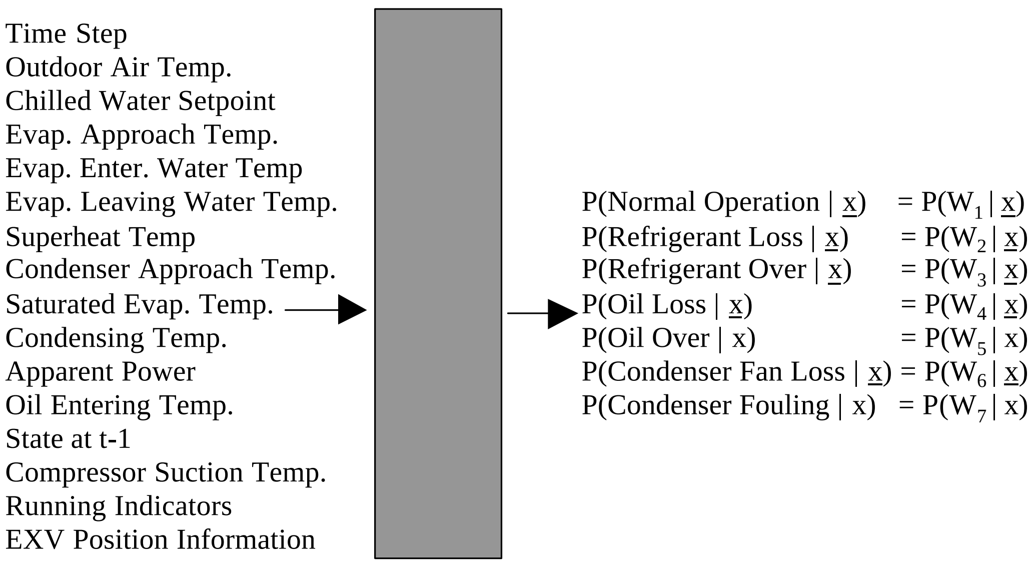

FIGURE 3. NN Classifier Input and Output Information

An important issue that is addressed early in the NN classifier development process involves the selection of specific type of NN classifier. Two common classes of NN designs include the radial basis function (RBF) and the multi-layered perceptron (MLP). RBF and MLP networks both provide effective methods of mapping between multi-dimensional input and output spaces. However, some important differences exist between the structure of each network class. The decision boundaries created by each class differ greatly. The MLP requires a training process that is highly non-linear with possible problems of local minima. This can lead to slower training than that required for the RBF network.

The network architecture of the two classes can also differ due to the relatively unlimited number of hidden layers that can be included within a MLP as compared with the RBF that typically includes only two layers. This difference can again affect the training time associated with the MLP. The final difference is in the actual training procedure used by each class. The MLP uses a global supervised training strategy that trains all of the network's parameters simultaneously. The RBF uses a two stage training process where the first stage is unsupervised. Again this difference can adversely affect the training time associated with the MLP. Refer to Bishop (1995) for a detailed discussion on each approach.

The pattern recognition problem addressed here included an input space that contained all levels of relevancy and significance as explored in Bailey (1998). Considering the nature of the input space and the potential difficulties associated with the radial basis function discussed above, the MLP approach was adopted.

Relevant Probability Theory

The following probability and decision theories were central to the design of the NN FDD system for vapor compression cycle equipment. For clarity, a list of definitions is included. The definitions have been adjusted to account for the problem's domain.

$\underline{x}$: vector of observed variables or chiller operating data.

C_i: current state of the equipment or chiller.

$p(C_i)$: prior or unconditional probability of the chiller state. This value expresses the probability of each state and is typically obtained by experts. The value represents an estimate of the probability that the chiller is in state C_i if no other information is given. The normal operating state for the chiller is given a much greater prior probability than any fault state.

$p(C_i|\underline{x})$: the posterior or conditional probability of the chiller state given the observed chiller data. The posterior probability is typically found using the theorem of probability called Bayes' Rule included in Eq.(4). The model that uses the posterior probability for classification is called a discriminative model because it classifies according to which chiller state has the highest posterior probability.

The classification process can be thought of as a set of discriminant functions $y_1(x),\ldots,$ $y_i(x)$ (Bishop 1995) with an input vector $\underline{x}$ assigned to class C_i if

$$y_i(\underline{x}) > y_j(\underline{x}) \text{ for all } j \neq i \tag{2}$$

where

$$y_i(\underline{x}) = P(C_i|\underline{x}). \tag{3}$$

Selecting the class C_i having the largest posterior probability minimizes the probability of misclassification. Posterior probability is typically found using the theorem of probability called Bayes' Rule. The model that uses the posterior probability for classification is called a discriminative model because it classifies according to which chiller state has the highest posterior probability.

The prior probability, $P(C)$, is the proportion of data within the training data set that is associated with a class label. In the chiller FDD model, prior probability correction is required. During the training and testing processes, the classifier produces an output or a posterior probability, $P(C_i|\underline{x})$ for each class. The classifier will predict posterior probabilities adequately for the training set, however, if the general population’s prior probabilities differ from that of the training set, the classifier will not perform adequately unless the following adjustment is made to the classifier's output.

$$Q(C_i \mid \underline{x}) = \hat{p}_i \left(\frac{Q(C_i)}{P(C_i)} \right) \qquad (4)$$

where

$\hat{p}_i$ = classifier's posterior probability of class i given observable $\underline{x}$

$P(C_i)$ = prior probability of class i in training set

$Q(C_i)$ = prior probability of class i data in testing set or general population

Therefore, the classifier's raw output must be corrected using a factor composed of the testing set prior probability for a specific class divided by the training set prior probability for that same class. For the chiller FDD classifier's design, the testing subset is not indicative of the general population. This is due to the large amount of fault data existing in the testing data set. Chiller data within the general population includes far less fault data. To account for the difference in data composition associated with the general population, prior probability, $Q(C_i)$, values for each class of chiller operation are estimated. These estimates are based on expert knowledge and chiller service records. Estimates of the general population's prior probability associated with each mode of operation are listed in Table 2.

TABLE 2.

Chiller Data Prior Probability Estimates for General Population

Chiller Mode of Operation	**Estimated Prior Probability for General Population**
normal	97.0%
refrigerant loss	0.5%
refrigerant overcharge	0.2%
oil loss	0.5%
oil overcharge	0.5%
air cooled condenser fan loss	0.05%
air cooled condenser fouling	0.75%
other	0.5%

Training, Testing and Validation Data

The data sets consist of chiller operating data collected during both normal and fault operation. The data used for classification purposes are collected using a chiller manufacturer's data acquisition system.

Data Separation. The hold out method is a technique commonly used in evaluating the performance of a classifier. This method is a simple approach to comparing different models by evaluating the error function of each model using a validation data-set that is independent of the training data-set. However, this method can lead to over-fitting of the validation set, therefore the performance of the selected model was confirmed by a third independent data-set called the test set. Therefore, data from each day’s testing were divided into three separate sets to be used for training, validation and testing.

During the data separation process, attempts are made to separate the data at natural time steps, such as when the chiller was cycled off due to low load or low evaporator discharge chilled water temperature. When the data separation phase was complete, the training data set generally includes data recorded during the first three hours of chiller operation. During this initial period, the chiller underwent the start-up process and two-thirds of the chiller loading process as prescribed by the load profile. The validation data set generally includes data recorded during the final one-third of chiller loading as prescribed by the load profile and a portion of the final steady state operating period. The testing data set generally includes data recorded during the final steady state operating period.

Data Preprocessing. Data preprocessing involves normalizing the input data and reducing its dimensionality. The normalizing process involved subtracting the mean ($\overline{x}_i$) of each feature and dividing by the standard deviation (s_i) as shown in Eq.(5). This simple linear normalization technique results in an input space with zero mean and unit standard deviation. No normalization was performed on the output space.

$$\tilde{x}_i = \frac{x_i - \bar{x}_i}{s_i} \tag{5}$$

Preprocessing of data can greatly enhance a classification system's performance by the inclusion of prior knowledge. The choice of a pre-processing technique is very critical because pre-processing is often the most important stage in the development of a classification model solution. One form of preprocessing is the reduction in dimensionality by removing unimportant data or features while maintaining as much relevant information as possible. This form of preprocessing is used in the current research through reducing the input space from fifty-two features to approximately thirty input features including evaporating, condensing, and superheat temperatures.

A network representation of the chiller FDD problem known as a belief network is developed by the author to assist in the complex process of reducing the dimensionality of the input space while maintaining as much relevant information as possible. The belief network diagram is included in Figure 4 and is created using expert domain knowledge. Belief networks are inference engines that efficiently and explicitly represent relevance relationships between variables. Therefore, a belief network is a tool that enhances the act of reasoning from factual knowledge or evidence. The network encodes relevancy using arrows that represent causal influences between neighboring nodes with the premise that consulting the neighborhood gives license to act. In other words, what does not exist locally in the neighborhood can be ignored. The belief network offers a clear semantic representation of the neighborhood associated with the chiller operating mode. Typically, belief networks play a vital role in uncertainty formalisms such as probability theory (Pearl 1988).

Using the results of the data visualization exercises (Bailey 1998), four key chiller variables were found to display significant effects due to operating mode. The four primary indicators of chiller health included current subcooling and superheating temperatures as well as condensing and evaporating refrigerant temperatures. In the belief network shown in Figure 4, the four variables are directly connected with the shaded box representing chiller operating state. If all four indicators of chiller health were available, the premise behind belief networks tells us that no other information is required to infer the chiller operating mode. However, one of the four values, subcooling temperature, is not available due to limitations of the chiller data acquisition system. Therefore, the three variables that directly affect and are directly affected by subcooling are shown in dashed boxes within Figure 4.

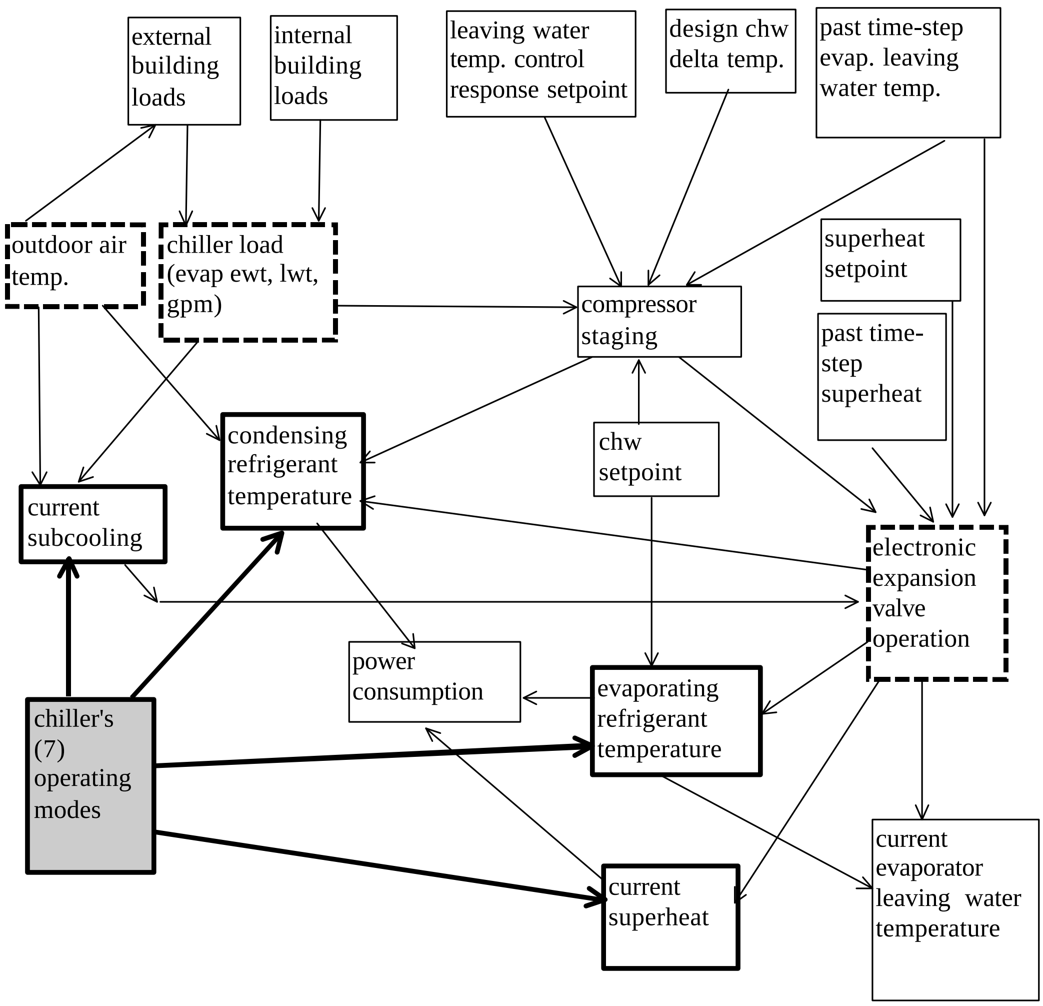

FIGURE 4. Belief Network Representation of Chiller FDD Problem

Training and Testing Process

First, a training data file was randomly selected from all training files while taking into account data file size. A training pattern was then randomly selected from the selected training data set. The training pattern included input data and associated target data. This type of training is known as supervised training since the desired output target value is known and used for training. Furthermore, the input pattern or vector was applied to the network input layer. Activation functions serve as a means of transforming the node's net input to an output. Activation functions studied for use within the NN FDD classifier include linear, sigmoid, tanh and softmax. It is important to note that the tanh activation function is used in hidden layers of single and multi-layered networks. However, in classification problems, when used in the output layer, the network's output can not be interpreted as a posterior probability due to the output transformation from negative to positive one.

After the training process is complete, a final weights file is created and the networks are tested recursively. A recursive network is one that uses sequential input data to predict values of an output. The term recursive implies that the output for any given time step is based not only on sequential input data but also on previous values of the output data. Such a technique is useful for systems such as chiller FDD where the current state depends on the recent state history. To train a network recursively, previous values of the known output are used as input values. That is, known output values at time steps k-1, k-2, k-3, etc. are used to predict the output at time step k. When testing a network recursively, the network must initially be "primed" with a known history of the output state. This initialized network is then used to predict the output at the next time step. The resulting output is then used as an input for the prediction at the next time step. The other, non-recursive inputs are treated in a normal fashion while the network outputs are recursed through this reintroduction into the network as inputs.

Testing is conducted using new data to determine the network's ability to map inputs to outputs. The testing program introduces the input data from the testing set to the network, which uses the inputs to predict an output value. The predicted value is then compared with the actual target value contained in the testing data set and an error was determined. The classification results of the trained networks were based on the misclassification rate as determined using Eq.(1).

Evaluation of Classifier Performance

A parametric study is conducted on the NN chiller FDD model using various arrangements of hidden layer architecture, activation functions, input histories and output recursions. Results are summarized in Table 3. For a detailed discussion on the parametric study, refer to Bailey (1998).

All of the input data are initially used for training and testing the various NN models until network12 was tested. Network12 was the first model tested without the data associated with the air-cooled condenser fan loss class. This action is taken after detailed analysis was conducted on the misclassification results obtained from configuration files network1 through network11. An exceedingly high misclassification rate is consistently associated with the air-cooled condenser fan loss class. The reasons for this can be attributed to the lack of subcooling data present within the training data set.

Network12 is the first model to remove the fan loss data from the data set for reasons discussed previously. The improvements to misclassification are obvious, especially when considering the testing data set as shown in Table 3. Network13 and network14 again included modified learning rates while network14 also used softmax activation functions on the output layer. The results for this model are not favorable and the softmax activation function is not used on any further models. The remaining models included varied input histories with identical output recursions, two hidden layers and sigmoid activation functions on both the hidden and output layers except for network19. Recall that activation functions serve as a means of transforming an input to an output within the NN model. Network19 incorporated three hidden layers and displayed poor results.

TABLE 3.

Summary of Results of Testing Networks with Training, Validation and Testing Data Sets

Configuration File	**Misclassification Rate Using Training Data Set**	**Misclassification Rate Using Validation Data Set**	**Misclassification Rate Using Test Data Set**
network1	7%	35%	69%
network2	66%	60%	82%
network3	11%	44%	76%
network4	17%	40%	44%
network5	13%	36%	84%
network8	3%	27%	58%
network9	1%	6%	28%
network10	0%	12%	56%
network11	0%	9%	36%
network12	0%	3%	20%
network13	0%	9%	47%
network14	28%	-	-
network15	31%	-	-
network16	2%	14%	40%
network17	0%	10%	28%
network19	5%	21%	42%

The best performing network results from network12, which is the first model trained without the fan loss data. This NN model includes 28 inputs (or 49 inputs including output recursion) with two hidden layers and seven output nodes. Sigmoid activation functions are used on the hidden and output layers. Impressively low misclassification rates of zero and three percent are associated with the training and validation data sets, respectively. Recall that the misclassification rate was determined using the following relationship.

$$Misclassification\ Rate = \frac{(number\ of\ misclassified\ samples)}{total\ number\ of\ testing\ patterns} \qquad \text{repeated}\,(1)$$

The misclassification rate associated with network12 based on the test data set is 20%. This increased error rate is attributed to the nature of the data in the testing data set that was primarily steady state operating data whereas the training and validation data sets included more transient operating data. The implications of the data separation procedure adopted in this research are an area of future research. Therefore, the results presented in Table 3 indicate that network architecture, activation function, input histories and output recursions each play a significant role in the performance abilities of the chiller FDD NN classifier.

Minimizing Risk

The classification decision used thus far has been based solely on maximum posterior probability. However, this may not be the most appropriate decision rule depending on the problem type. Perhaps a class dependent cost should be associated with misclassifying a sample. A good example to illustrate this point is a classifier that distinguishes normal tissue from cancer tissue. The cost associated with classifying a normal tissue as cancerous is considerably less than that associated with misclassifying a cancerous tissue as normal.

A simple, straightforward approach to account for these effects is through the creation of a cost matrix. The cost matrix includes elements l_{ij} that stipulates the fine or penalty associated with assigning a pattern to class C_i when it belongs to class C_j. The values within a cost matrix intended for the chiller FDD problem would be provided by experts and would therefore be subjective in nature.

The conditional risk associated with the action (a_i) of assigning a pattern to class C_i when it in fact belongs to class is included in Eq.(6).

$$R(\boldsymbol{a}_i \mid \underline{x}) \equiv \sum_{j=1}^{k} \boldsymbol{l}(\boldsymbol{a}_i \mid C_j) P(C_j \mid \underline{x}) \tag{6}$$

where :

$R(\boldsymbol{a}_i \mid \underline{x})$ = the risk associated with assigning a pattern to class C_i

k = the total number of classes

$\boldsymbol{l}(\boldsymbol{a}_i \mid C_j)$ = the cost associated with assigning a patern to class C_i when it belongs to class C_j

$P(C_j \mid \underline{x})$ = the posterior probability of class C_j given the observeable $\underline{x}$

The new decision rule involves selecting the action associated with the lowest risk as shown in Eq.(7) where action $\boldsymbol{a}_i$ would be chosen.

$$R(\boldsymbol{a}_i \mid \underline{x}) < R(\boldsymbol{a}_j \mid \underline{x}) \qquad \text{for all } i \neq j \tag{7}$$

The incorporation of risk analysis through the creation of a cost matrix can yield a FDD system that produces very low rates of false alarms as shown by Dodier et al. (1998). Because this implementation is quite straightforward and has been demonstrated in previous research, it was not formally implemented in the FDD classifier developed here. However, the method's worth and potential benefits should be considered in future research efforts in this area.

Another possible means of reducing both the rate of false alarms and the misclassification rate is through the adoption of rejection thresholds. For most of the misclassified samples in this research, the largest posterior probability was relatively small due to overlap within the input space. In such cases, it may be wiser to not make a decision at all unless the maximum posterior probability exceeds a predetermined threshold.

Therefore, if the posterior probability fell below the threshold, the classifier's output would be the unknown state. Thereby, a seven-state classifier would be converted to an eight-state classifier that incorporates the unknown mode of operation. As with the cost matrix discussed above, the rejection threshold would be determined by experts and would therefore be subjective. Because this implementation is also quite straightforward and has been demonstrated in previous research, it was not formally implemented into the chiller FDD classifier. However, like the cost matrix, the addition of rejection thresholds could be extremely useful in future research efforts in this area.

CONCLUSIONS

This paper discusses the process involved in designing and testing the viability of a NN based FDD system for chillers. The fault detection and diagnostic tool performs the task of analyzing archived chiller operating data and detecting faults through recognizing patterns within the data. The NN program is trained on chiller operating data collected during normal and fault conditions in a full-scale HVAC laboratory. Faults include the improper charging of refrigerant and oil as well as air-cooled condenser fouling and fan loss. The final, trained FDD tool is capable of correctly predicting the current state of chiller operation with a misclassification rate of three percent when presented with never before seen data sets. Future and on-going research related to this paper includes testing the FDD model's applicability to other chiller types, sizes and on site arrangements.

ACKNOWLEDGEMENTS

I extend my gratitude to Erik Jeannette and Loretta Cirricione for their invaluable assistance during the extensive laboratory testing process required completing this research. In addition, the discussions with several Trane employees were useful. The National Science Foundation and The Trane Company supported this work.

REFERENCES

Bailey, M.B. [1998]. System performance characteristics of a helical rotary screw air-cooled chiller operating over a range of refrigerant charge conditions. ASHRAE Transactions, Vol. 105, Part 2.

Bailey, M.B. [1998]. The design and viability of a probabilistic fault detection and diagnosis method for vapor compression cycle equipment. Ph.D. Dissertation, Department of Civil and Architectural Engineering, University of Colorado.

Bailey, M.B., Kreider, J.F. [2000]. Experimental methodology utilized for chiller fault simulations. Proceedings from the International Instrumentation Symposium. Bellevue, WA, April 30- May 4.

Barron, A.R. [1993]. Universal approximation bounds for superpositioning of a sigmoidal function. IEEE Transactions on Information Theory, Vol. 39, No. 3, pp. 930-945.

Bishop, C.M. [1995]. Neural networks for pattern recognition, Clarendon Press, Oxford, England.

Bridle, J.S. [1990]. Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition. Neurocomputing: Algorithms, Architectures and Applications, F.F. Soulie', J. Herault, eds., pp. 227-236. New York: Springer-Verlag.

Dodier, R., Curtiss, P. Curtiss, Kreider, J.F. [1998]. Small scale on-line diagnostics for an HVAC system. ASHRAE Transactions, Vol. 104, Part 1.

Kreider, J.F., Curtiss, P.C., Massie, D., Jeannette, E. [1999]. A commercial-scale university HVAC laboratory. ASHRAE Transactions, CH-99-13-4.

Massie, D., Bailey, M.B. [2000]. Optimization and fault evaluation of a building's cooling plant using neural networks. Proceedings from ECOS 2000 International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Aspects of Energy and Process Systems, University of Twente, Enschede, The Netherlands, July 5-7.

Pearl, J. [1988]. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers, Inc. San Mateo, California.