ACOUSTIC CHARACTERIZATION OF PROSTHETIC HEART VALVES

G. H. Thomas, J. V. Candy, D. Perkins, R. D. Huber, and M. Axelrod
Lawrence Livermore National Laboratory
P. O. Box 808 L-333
Livermore, CA 94550

INTRODUCTION

Prosthetic heart valves are a blessing for people with defective heart valves. One type of mechanical heart valve that was manufactured between 1979 and 1985 has been implanted in approximately 86,000 people. Between 500 and 600 of these valves are known to have failed. The failure occurs when a thin wire strut breaks free, see Figure 1. The strut has two legs that are attached to the main body of the heart valve. Typically one of the legs breaks first, leaving the other leg intact and the heart valve still functioning. This condition is called a single leg separation. A technique that analyzes the sound generated by the heart valve was developed to detect this single leg separation. Acoustic data was acquired from implanted heart valves prior to their being explanted. These signals were processed and distinguishing characteristics have been identified that correlated the condition of the heart valve strut with its acoustic signature. An automated classification algorithm was developed and trained to predict the heart valve's condition.

The classification protocol is shown in Figure 2. The first task is to extract the closing and opening heart valve sounds from the measured acoustic recordings. Then the signals are processed to extract pertinent features which distinguish between intact (INT)

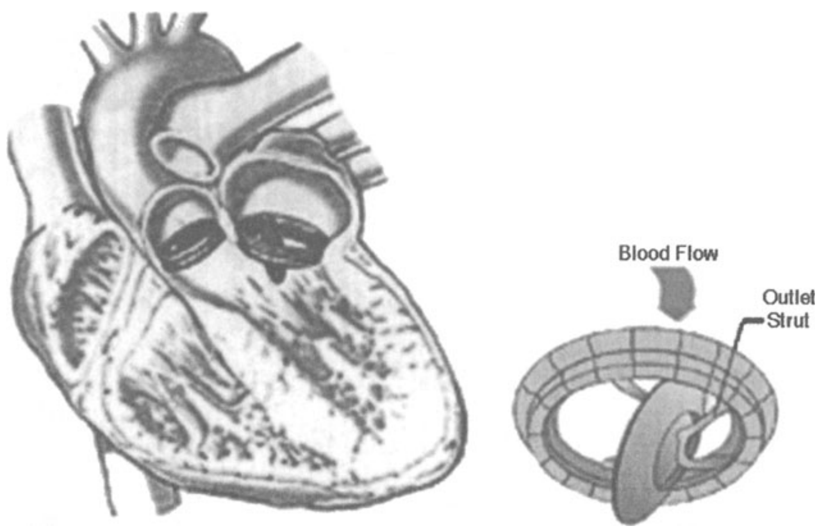

Figure 1. Prosthetic heart valve. A single leg separation occurs in the outlet strut.

Review of Progress in Quantitative Nondestructive Evaluation, Vol. 18
Edited by Thompson and Chimenti, Kluwer Academic/Plenum Publishers, 1999

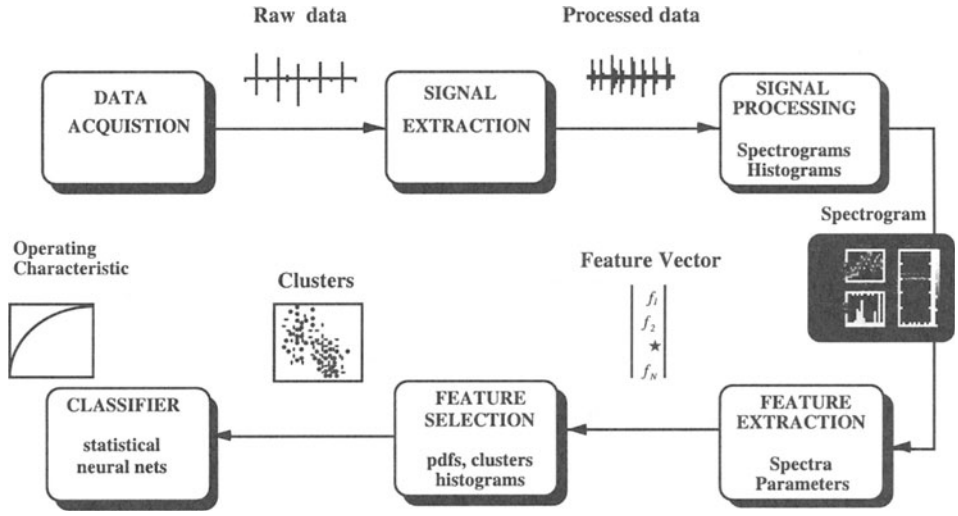

Figure 2. Classification protocol used for heart valve condition determination

struts or single leg separated (SLS) struts. We, therefore, classify the condition of the valve under test as INT or SLS. As shown in Figure 2, once the features are selected, they are used to calculate the required decision metric which is compared to a prescribed threshold to classify the condition of the heart valve under test.

Signal Processing

Before a heart valve classifier can be developed, it is necessary to minimize the adverse distortions (i.e. noise) in the acoustic signals [1]. Signal processing algorithms were written to extract acceptable signals from the acoustic data, minimize the noise influence, and transform the individual signals so that features could be selected.

Feature Extraction

The signal processing tool provides a means of interactively handling heart valve acoustic data. Since the human brain is an excellent pattern recognizer, it is important to visualize the data and discern the best processing techniques. The acoustic heart valve data was processed to recognize trends and correlations between SLS (single leg separation) and INT (intact) valves. The features investigated included: frequency spectra, coefficients from spectral models (including lattice parameters, predictor parameters, and correlation coefficients), first-order statistical features, second-order texture features, and other representations of the data [2,3,4,5,6,7,8,9,10].

Feature Selection

We implemented automatic and manual feature selection procedures. Algorithms automatically searched through the set of features and ranked them by order of importance. One algorithm ranked the features one by one in order of a statistical measure of distance between cluster centers in feature space. Another algorithm produced a list of the optimal set of features (given a number of features to choose a priori). After using feature selection algorithms, we performed a manual inspection of the one- and two-dimensional cluster plots in feature space to reduce the feature set, to gain physical insight and to allow the insertion of valuable human judgment into the process. We choose the number of features according to the method that says that the number of independent training samples (feature vectors) should be greater than or equal to approximately 5 times the number of features contained in a feature vector. Thus the number of features applied the classification is limited by the number of training samples (i.e. heart valves) available.

Classification

The goal of the classification task is to determine the state of an unknown heart valve (intact or single leg separation) from characteristics or features of the acoustic signal generated by the functioning valve. A wide variety of classifiers is available, including: nearest neighbor [6], linear discriminants [6], back-propagation neural networks [11] and probabilistic neural networks [12]. These classifiers infer the valve condition by comparing the features extracted from its sounds with those features extracted from sounds of known valves.

Neural networks are parallel, distributed information processing structures consisting of interconnected processing elements. The processing elements can possess a local memory and can carry out localized information processing operations. Most neural network algorithms are adaptive systems which use heuristic approaches to discover underlying class statistics. The heuristic approaches usually involve making many small changes to the system parameters that gradually improve system performance

The Probabilistic Neural Network [12] was selected for classifying the heart valve condition. The Probabilistic Neural Network has interconnection structures based upon statistical principles. Instead of minimizing a performance criterion such as mean-square error between known and estimated system outputs, the Probabilistic Neural Network estimates probability density functions of the random variables involved in the process being modeled. It has the unique property that under certain easily met conditions, the decision surface estimated by the Probabilistic Neural Network asymptotically approaches the Bayes optimal surface, as the sample size increases.

Our feature based classification protocol was based on two sets of acoustic data. The first set of data trains the feature extraction and pattern recognition system. This training data set is made up of acoustic signals from each class of heart valve. Classification algorithms are produced by the computer based on the expected results from the training data. The confidence in the classification results becomes greater as the size of the training set increases. An iterative process is invoked on the training data until an acceptable algorithm is derived from a combination of the best acoustic signal features and the best pattern recognition algorithm. The performance of the classification algorithm was measured by classifying the acoustic signals from a set of test data. Classifying the condition of the heart valves represented in this test set determines the specificity and sensitivity of the classifier. Increased confidence in the heart valve classifier is realized as the sets of training and test data grows. The result of the signal processing, feature extraction and pattern recognition is a complete system to classify heart valve outlet strut condition from its acoustic signature.

RESULTS

Training

The first step in developing a classification algorithm is to train using known valves for each class. This is called supervised learning. In this study we had a total of 23 valves (13 INT, 10 SLS) to train the classifier. After carefully pre-processing the sounds for each training valve, those beats qualifying as acceptable were extracted and incorporated into the training set for supervised classification. These supervised classifiers can only perform as well as the training set data input during the supervised training. This requires that the training data be a valid representation of the possible valve conditions to be classified. Once the training set has been extracted and processed, a "self-consistency" check is performed on this set called cross-validation or more colloquially hold-one-out. This calculation involves: (1) removing each valve in the training set--one at a time (hold-one-out); and (2) classifying the (already known) condition of the valve held out to assure it is still represented by the remaining valves. Theoretically, the classifier should be able to classify each valve in the training set correctly validating its performance.

Once the training set of valves is constructed, the performance of the classifier can be characterized by its operating characteristic or equivalently receiver operating

Table 1. Training results for heart valve condition classifier.

No. of Valves:	23 (13 INT, 10 SLS)
Correct Classification (x 100):	100 %
Sensitivity (x 100):	100 %
Specificity (x 100):	100 %
Spectral Features (center freq., bw):	(13.6 kHz, 469 Hz)
	(16.7 kHz, 469 Hz)
	(20.0 kHz, 94 Hz)

characteristic (ROC) curve. The ROC is a measure of sensitivity (probability of detection) versus (1-specificity) or probability of false alarm and can only be obtained from the "training set" data, since the condition of the valves must be known a priori to construct the curve. Once the ROC is determined, an operating point for the classifier must be selected from the curve based on the desired sensitivity/specificity. A threshold value corresponding to this operating point is then invoked for the blind set classification.

Our choice of the feature type for this study is based on the vibrational response of the heart valve. The vibration response is characterized by spectral bands of acoustic energy, thus the valve sounds are transformed into the frequency domain and decomposed into spectral bands. The spectral bands are then automatically sorted by the computer and ranked from the most sensitive to the least sensitive to the valve condition. This procedure is the feature selection operation, see Figure 2. The three spectral bands automatically chosen by the feature selector were (center frequency, bandwidth): (13.6 kHz, 469 Hz), (16.7 kHz, 469 Hz), and (20 kHz, 94 Hz). These features were incorporated into the probabilistic neural network algorithm to derive a classification predictor.

The results of our training effort as shown in Table 1 were excellent. Note that not only were the valves classified, but almost all sounds were perfectly classified giving posterior probabilities that were Prob(SLS/SLS)=1.0 and Prob(INT/SLS)=0.0 with extremely high confidence in the posterior probability estimates (negligible intervals).

Blind Test

Once the classifier was developed from the training data set, it was tested on a blind data set supplied by the sponsor. After carefully processing the data, the blind test sounds were extracted and classified using the spectral features and the corresponding threshold selected during training, see Table 1. The results are summarized in Table 2. The classification results for each of the 50 valves are shown in Figure 3 along with their corresponding confidence interval. After the classification was performed, we obtained the true valve class (ground truth) and annotated each correctly classified valve by a "diamond" and each misclassified valve by a "circle" on the plot. Also shown is the valve's predicted confidence interval. An arbitrary decision threshold was established at P(SLS/X) of .5. For the blind test, the classifier achieved a sensitivity of 80% and a specificity of 66%. A more meaningful overall measure is the probability of correct classification (x 100) which was 72%.

Table 2. Test results for heart valve condition classifier.

No. of Valves:	50 (30 INT, 20 SLS)
Correct Classification (x 100):	72 %
Sensitivity (x 100):	80 %
Specificity (x 100):	66 %

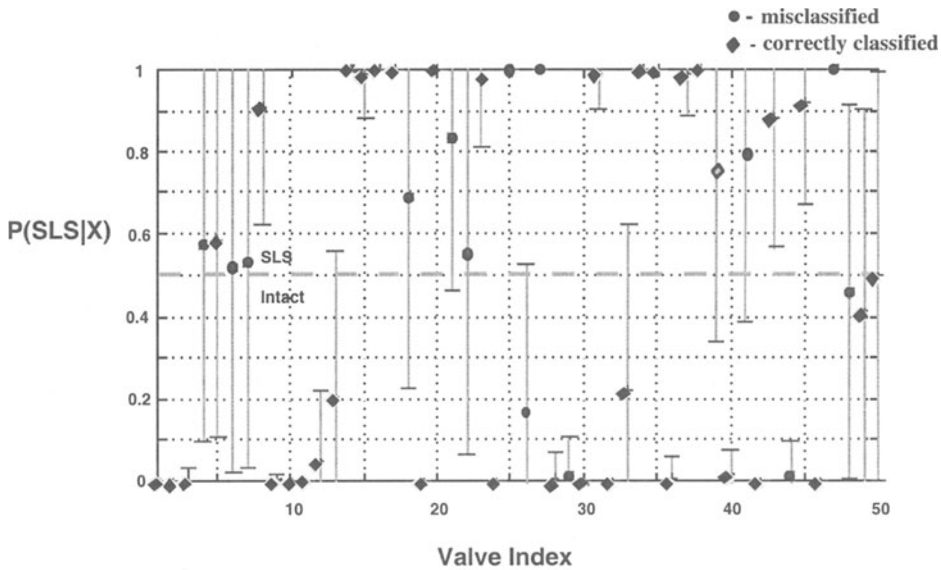

Figure 3. Results of classifying blind set of heart valve acoustic data.

Table 3. Test results for heart valve condition classifier after removing "no calls".

No. of Valves:	35 (19 INT, 16 SLS)
Correct Classification (x 100):	85.7 %
Sensitivity (x 100):	87.5 %
Specificity (x 100):	84.2 %

In Figure 3, the predicted estimates that cross the 0.5 decision threshold may be defined as the "no-calls". It is interesting to note that if we exclude the no-calls due to high uncertainty in the posterior probability estimate, then we eliminate 15 sessions and our probability of correct classification increased to 85.7% with a corresponding sensitivity of 87.5% and specificity of 84.2%, see Table 3.

CONCLUSION

We demonstrated that the acoustic signatures of the functioning heart valves provide information as to the condition of the outlet strut. We developed advanced signal processing techniques to extract the acoustic signals and transform those signals into the frequency domain to provide features for classification of the heart valve. Our classification procedures have produced excellent results on a limited number of heart valves. We will increase the confidence in the predictive capabilities of the heart valve classification algorithms as we process additional data.

ACKNOWLEDGMENTS

Work performed under auspices of the U. S. Department of Energy by the Lawrence Livermore National Laboratory under contract No. W-7405-ENG-48.

REFERENCES

1. Candy, J. and Barnes, F., Heart valve processing: a feasibility study. Lawrence Livermore National Laboratory Report. (1991) UCRL-ID-107630.

2. del Grande, N.K., Clark, G.A., Durbin, P.F., Fields, D.J., Hernandez, J.E., and Sherwood, R.J. Buried Object Remote Detection Technology for Law Enforcement. SPIE Orlando '91 Symposium, Orlando, Florida, April 1-5, 1991.
3. Ballard, Dana H. and Brown, Christopher M. Computer Vision. Englewood, NJ: Prentice-Hall; 1982.
4. Jain, A.K. Fundamentals of Digital Image Processing. Englewood, NJ: Prentice-Hall; 1989.
5. Weschler, H. Computational Vision. San Diego, CA: Academic Press; 1990.
6. Duda, R.O. and Hart, P.E. Pattern Classification and Scene Analysis. New York, NY: Wiley; 1973.
7. Kohonen, T. Self-Organization and Associative Memory. New York, NY: Springer-Verlag; 1989.
8. Mar, D. Vision. New York, NY. W. H. Freeman and Co; 1982.
9. Pratt, W.K. Digital Image Processing. New York, NY: Wiley; 1978.
10. Rosenfeld, A. Computer Vision: Basic Principles. Proceedings of the IEEE. 1988; 76:863-8.
11. Rumelhart, D. E. and McClelland, J.L., "Parallel Distributed Processing: Explorations in the Microstructure of Cognition" Vol. 1: Foundations, MIT Press, 1986.
12. Specht, D. F., "Probabilistic Neural Networks," Vol. 3, Neural Networks, 1990.