

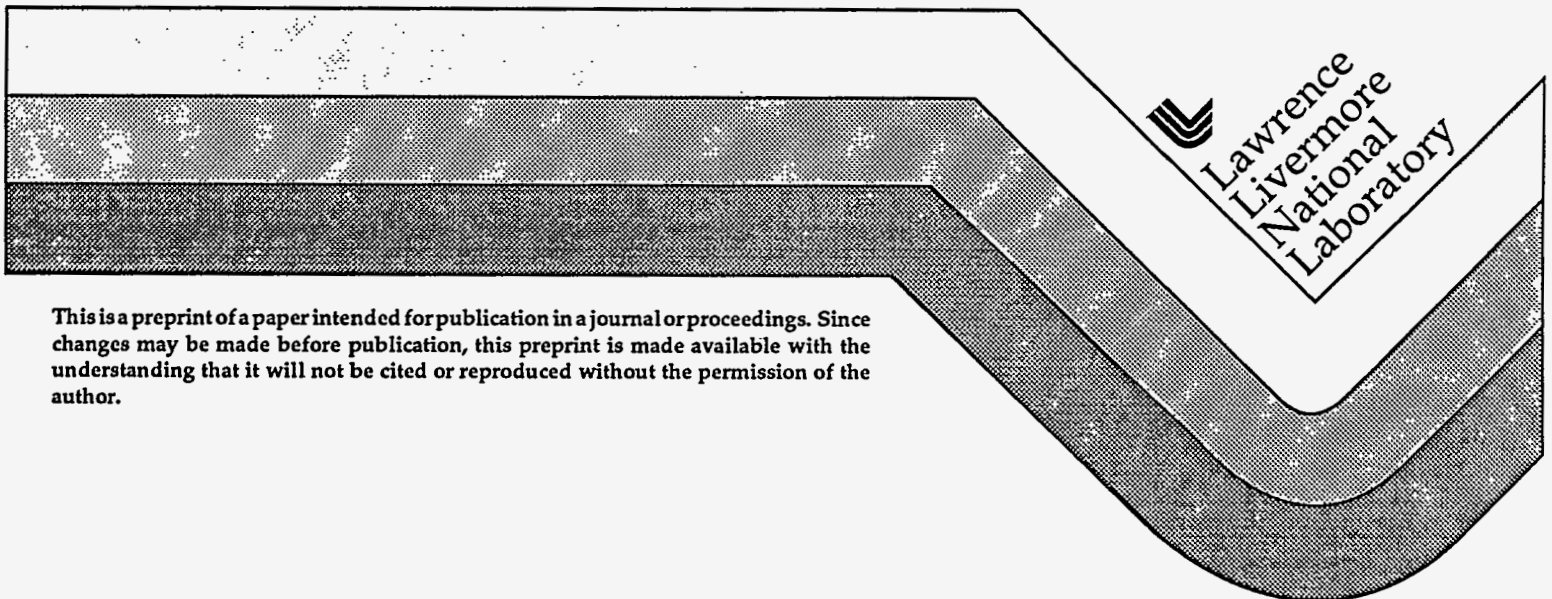
UCRL-JC-117081
PREPRINT

**Signal Processing and Classification of Acoustic
Signatures from Björk-Shiley Convexo-Concave
Heart Valves**

**Graham H. Thomas, Gregory A. Clark, Susan L. Crawford,
Michael R. Buhl, and Nick Borate**

**This paper was prepared for submittal to the
Shiley Heart Valve Research Center Acoustic Research Workshop
Livermore, California
May 2-3, 1994**

April 1994



This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

jo

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Signal Processing and Classification of Acoustic Signatures from Björk-Shiley Convexo-Concave Heart Valves

Graham H. Thomas, Gregory A. Clark, Susan L. Crawford,
Michael R. Buhl, and Nick Borate

Lawrence Livermore National Laboratory
Livermore California, 94550

Abstract

The prosthetic heart valve has improved the length and quality of life for people with serious heart conditions. Even though the designs are extremely reliable, the valves are mechanical and operating continuously over a long period, therefore, structural failures can occur. Lawrence Livermore National Laboratory is performing the research and development necessary to construct an algorithm capable of non-invasively classifying the condition of the outlet struts of implanted Björk-Shiley Convexo-Concave (BSCC) heart valves. This technique will analyze acoustic signals that are caused by the heart valve's disc striking the outlet strut when the valve opens. Since the disc is activating the outlet strut directly, the condition of the strut will have the greatest influence on the radiated acoustic signals. These acoustic signals are recorded in vivo from functioning heart valves. Our approach is to apply the signal processing necessary to extract pertinent information from the acoustic data and develop BSCC heart valve classification algorithms based on features of the enhanced acoustic signals. We have assembled the hardware and software needed to extract opening sounds from acoustic data sets, perform quality control on the signals, reduce the size of the data set by various transformations, extract features from the transformed signals, select the optimal features for classification, and build classification algorithms to determine the condition of the valve from its acoustic signature. We have identified several parameters (features) of the opening signals which we process with neural network classifiers to predict heart valve condition. These algorithms have proven effective for classifying a limited number of explanted heart valves. We continue to expand our heart valve data set and refine the classification techniques.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Introduction

The goal of this research project is to develop a technique to non-invasively monitor the acoustic signals produced by BSCC heart valves to determine the structural condition of the valves. Since intact valves radiate different acoustic signals than single leaf separation (SLS) valves, a detailed analysis of the acoustic radiation will indicate the condition of the valve. Before a heart valve classifier can be developed, it is necessary to minimize the adverse distortions (i.e. noise) in the acoustic signals (Candy, 1994). Once the

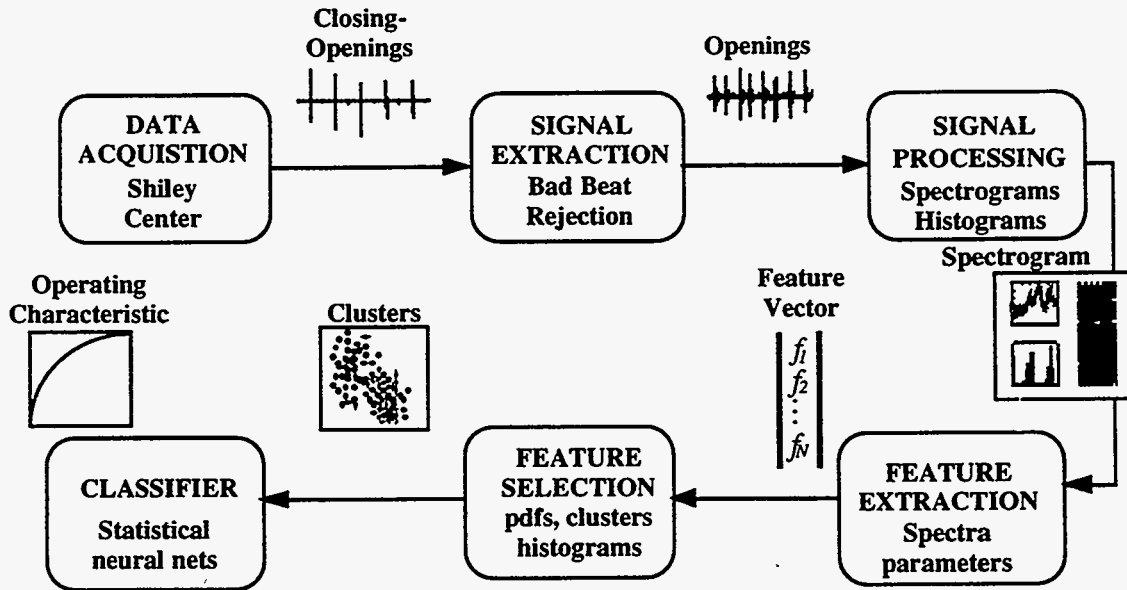


Figure 1. Heart valve classification protocol.

signal processing techniques have been developed and the characteristics of the acoustic signals that identify intact and SLS heart valves are determined, we will build a software protocol to automatically classify the BSCC heart valve condition. A statistically significant amount of acoustic data will be processed to confirm the required performance of the classification techniques.

Early studies conducted by LLNL have proved that it is possible to determine valve condition from the acoustic signal radiating from functioning heart valves in sheep and in people. The program first conducted on ovine acoustic data provided valuable insight for our recent studies on clinical data (Mullenhoff, 1993, Crawford, 1993, Buhl, 1993). There are differences in the acoustic signals generated by BSCC heart valves with SLS and intact outlet struts. Through careful analysis, one may non-invasively determine the

condition of an implanted BSCC valve by interpreting these sounds and resonant frequencies.

The classification process consists of two main steps: the training phase and the testing phase. The goal of the training phase is to learn as much as possible about the process from a known set of data (training data). This data was obtained from explanted valves, and therefore their condition was positively known. The goal of the testing phase is to use the information learned in the training phase to classify a "blinded" (unknown to the processor) data set.

Feature Extraction

Features can be selected from a priori knowledge (i.e., pulse duplicator studies), from computer modeling of the acoustic response of heart valves, and from human interpretation. The signal processing prototyping tool developed by LLNL provides a means of interactively manipulating 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 will be processed to recognize trends and correlations between SLS and intact valves. The features to be investigated will include the following: frequency spectra, coefficients from spectral models (including lattice parameters, predictor parameters, correlation coefficients, and others), first-order statistical features, texture features (second-order), and other representations of the data (del Grande, 1991, Ballard, 1982, Jain, 1989, Weschler, 1990, Duda, 1973, Kohonen, 1989, Marr, 1982, Pratt, 1978, Rosenfeld, 1988). It is anticipated that LLNL will be able to use these feature types both individually and in concert to maximize classification performance. LLNL will work with SHVRC and other collaborators to identify viable features of the acoustic signals.

Feature Selection

We use both automatic and manual feature selection procedures. We have advanced algorithms for automatically searching through the set of features and ranking them in order of importance. For example, some algorithms rank the features one by one in order of a statistical measure of distance between cluster centers in feature space. Other algorithms produce a list of the optimal set of features (given a number of features to choose a priori). After using feature selection algorithms, we generally perform a manual inspection of the one- and two-dimensional cluster plots in feature space to further 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 we can use is limited by the number of training samples (valves) available.

Classification

The goal of the classification task is to determine the state of an unknown heart valve (intact or SLS) from characteristics or features of the acoustic signal generated by the functioning valve. A wide variety of classifiers is available, including: nearest neighbor (Duda et. al., 1973), linear discriminants (Duda, et. al., 1973), back-propagation neural networks (Rumelhart et. al., 1986) and probabilistic neural networks (Specht, 1990). These classifiers infer the valve condition by comparing the features extracted from its opening sounds with those features extracted from openings 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. Neural networks have been applied to many pattern classification, communication, and control problems. 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. An example of this type of neural network is the back propagation neural network (BPNN) (Rummelhart et. al., 1986).

Recently, a different neural network paradigm called the Probabilistic Neural Network (Specht, 1990). The PNN, while having interconnection structures similar to the BPNN are based upon statistical principles. Instead of minimizing a performance criterion such as mean-square error between known and estimated system outputs (as the BPNN does), The PNN 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 PNN asymptotically approaches the Bayes optimal surface, as the sample size increases.

The Probabilistic Neural Network (PNN) has several advantages over the BPNN, making it a very effective classifier for most applications. Among these are the following: (1) it

requires much less training time (many orders of magnitude) for many problems. (2) unlike the BPNN, the PNN cannot converge to poor solutions corresponding to local minima of the error criterion, because it is not an adaptive technique. It is a nonparametric estimator of the probability density functions of the random variables representing the classes of interest, given the observed samples. (3) it has better generalization capability than the PNN. (4) it learns in just one pass through the data and can generalize from examples as soon as they are stored. (5) unlike the BPNN, the PNN often performs very well when the set of training patterns is sparse, with a smooth transition from one sample point to another. The main drawback of the PNN is that all training samples must be stored and used in classifying new patterns. However, this is not a serious problem, because once the decision boundary has been formed, it can be approximated by polynomial hypersurfaces of sufficiently large degree (Specht, 1990) and then the coefficients of these polynomials instead of the actual samples can be stored. Furthermore, these stored coefficients are easily updated as new samples arrive. Because memory is inexpensive, storage is not a problem for moderate size applications, such as the heart valve analysis problem.

LLNL feature based classification protocol is 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 tested.

Classification algorithms are calculated 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 is 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. As the heart valve classification algorithm builds on growing sets of training and test data, more confidence is realized in its classification ability. The result of the signal processing, feature extraction and pattern recognition will be an complete system to classify heart valve outlet strut condition from its acoustic signature.

Our approach is to begin the study by limiting the signal variables to just those caused by the valve. In this phase nearest neighbors and linear discriminants are adequate for SLS detection. Then as the classification requirements become more complicated (i.e.

different size valves, patients with double valves, and various heart conditions), advanced algorithms such as neural networks will be implemented to detect SLS. Because both neural network algorithms make no prior assumptions about what features are relevant to the classification, the networks produce non-intuitive boundaries which would be very difficult for a human to realize. Since these decision boundaries are developed automatically, they will recognize the different valve size as well as the condition of the valve.

Since the heart valve classifiers, both classical and adaptive, require large amounts of independent data sets (different heart valves) for the design and performance analysis, we propose to develop a heart valve synthesizer. This beat synthesizer will generate the necessary number of independent acoustic data sets. The synthesizer will be based on actual data extracted from the clinical studies. This synthesizer will provide us with an enlarged ensemble of acoustic data that represents realistic signals to improve the training of the classifiers. Also modifying the synthetic data to represent acoustic signals outside the range of actual data, helps determine the robustness of the classifier.

Results

Our study of clinical acoustic data has demonstrated the feasibility of extending the techniques developed for ovine heart valves to human heart valves. We have conducted an organized protocol for classifying heart valve condition from its acoustic signature, see Figure 1. The data acquisition and signal processing efforts are described in a paper by Candy and Jones (Candy, 1994). The classification portion of this work is presented in this paper.

The first study of the clinical data demonstrated that there is information in the opening acoustic signal which indicates the heart valve condition. A set of data from nine confirmed heart valves was processed as the training data for a Fisher Linear Discriminant classifier (Duda, 1973). Each opening signal was considered a separate data due to the limited amount of data. The feature source was spectrogram data files calculated with a Lattice spectral estimation technique after high pass filtering and beat rejection. The SLS data was combined into one file and the intact data was combined into another file. These two data sets were processed with feature extraction and pattern recognition algorithms which had proven valuable for classifying the ovine data. The

features selected represented energy in 1 KHz bands over the 10 to 24 KHz range, see Table 1.

Table 1. Features selected for phase one clinical data classification, represent energy in 1 KHz bands over the 10 to 24 KHz range.

<u>Feature Number</u>	<u>Definition</u>
1	Energy in the frequency band from 10 to 11 KHz
2	Energy in the frequency band from 11 to 12 KHz
3	Energy in the frequency band from 12 to 13 KHz
4	Energy in the frequency band from 13 to 14 KHz
5	Energy in the frequency band from 14 to 15 KHz
6	Energy in the frequency band from 15 to 16 KHz
7	Energy in the frequency band from 16 to 17 KHz
8	Energy in the frequency band from 17 to 18 KHz
9	Energy in the frequency band from 18 to 19 KHz
10	Energy in the frequency band from 19 to 20 KHz
11	Energy in the frequency band from 20 to 21 KHz
12	Energy in the frequency band from 21 to 22 KHz
13	Energy in the frequency band from 22 to 23 KHz
14	Energy in the frequency band from 23 to 24 KHz

Note: High lited features had the best correlation with heart valve condition in this phase I study.

The average feature values for intact and SLS data is plotted in Figure 2. The error bars represent one standard deviation. The average SLS and intact spectrum are plotted in Figure 3. Highlighted in Figure 3 are the two bands of energy which performed best in discriminating between intact and SLS valves. From the spectra it appears that the energy in the 14 to 15 KHz range (Feature 4) would be a good discriminator but looking at the average and standard deviation values for this corresponding feature, one notes a large standard deviation associated with the SLS data for this feature (see Figure 2). Thus this feature was not included in the classification procedure. In processing the data during the training phase, the features are analyzed individually and then in

combinations. Feature 7, 11 and 12 performed the best on an individual basis. The results are displayed in Table 2:

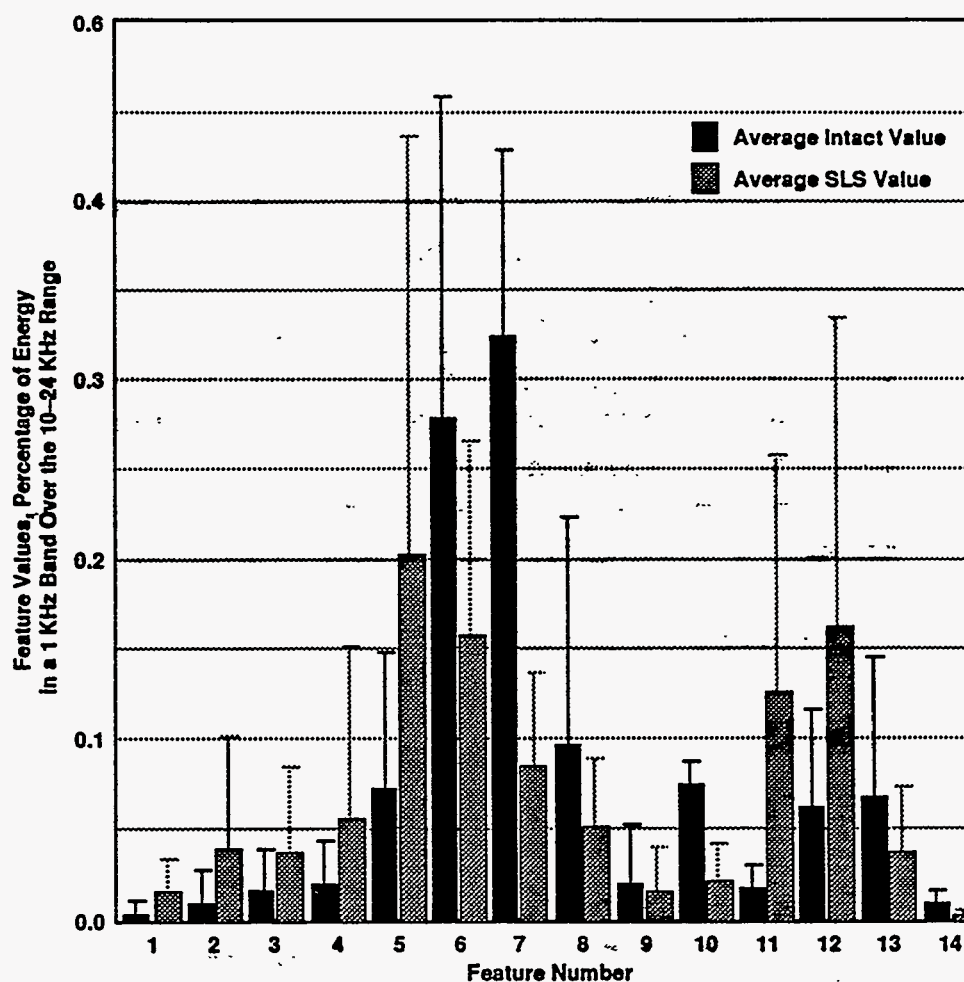


Figure 2. Feature values averaged from 923 intact and 1782 SLS heart beats (4 intact valves and 5 SLS valves). Features 7, 11, and 12 have discrimination value.

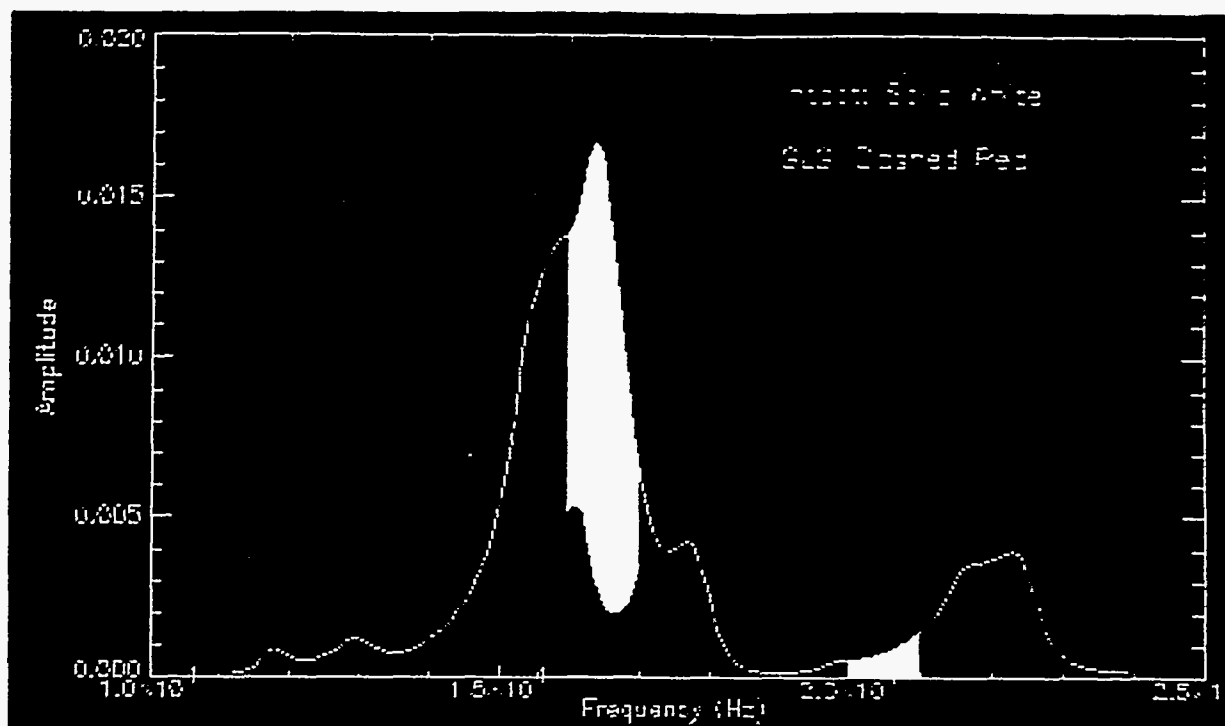


Figure 3. Frequency analysis of the opening sounds displays a distinction between SLS and intact valves. Note: Difference in energy content between SLS and intact valve sounds in the 16–17 KHz and 20–21 KHz bands.

Table 2. Classification results for phase one clinical study.

Feature Number	Intact Data		SLS Data	
	<u>Classified I</u>	<u>Classified SLS</u>	<u>Classified I</u>	<u>Classified SLS</u>
7	828 (89 %)	95	36	1746 (98 %)
11	921	2	969	813
12	761	162	1033	749
7 & 11	849 (92 %)	74	36	1746 (98 %)

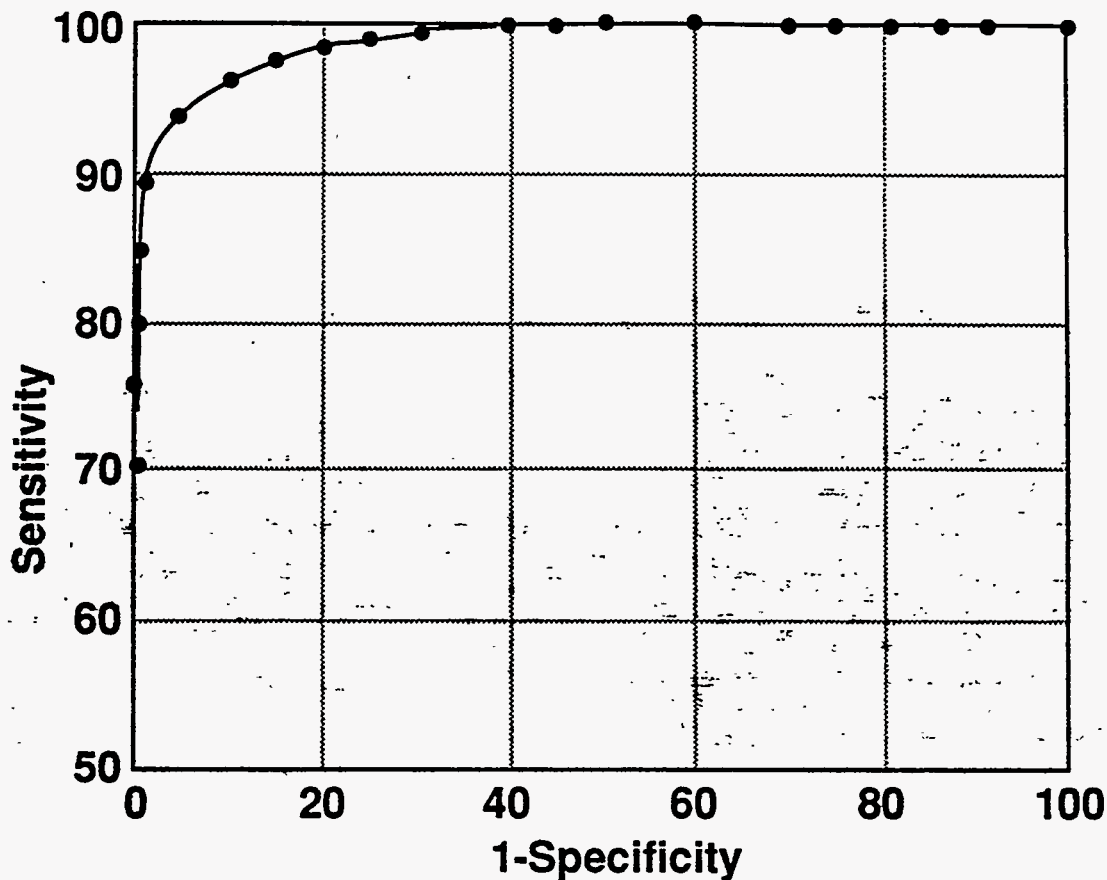


Figure 4. ROC display of results for a Fisher Linear Discriminant classifier for feature 7, (energy in the 16–17 KHz region)

From these results one notes that the intact data is correctly classified by all three features while the SLS data is correctly classified only by feature 7. Both features 11 and 12 misclassify the SLS data. An ROC curve was generated for feature 7 and is shown in Figure 4. At a threshold of 0.5 the sensitivity, (probability of detection) is 89.5% and the (probability of false alarm), $1 - \text{specificity}$, is 2%. When feature 11 and 12 are used together the performance is improved slightly. At a 0.5 threshold the sensitivity is 92.0% and the false alarm is 2.0%. The corresponding ROC curve is shown in Figure 5, (page 11). Two space plots can show the clustering or separation of the SLS and intact data based on the feature values. The best two space plot for this data set is obtained by plotting feature 7 data against feature 11 data and is shown in Figure 6, (page 12). The green stars represent the centers of the intact and SLS data. The line drawn is the perpendicular bisector which aides the visualization of natural clustering of the data.

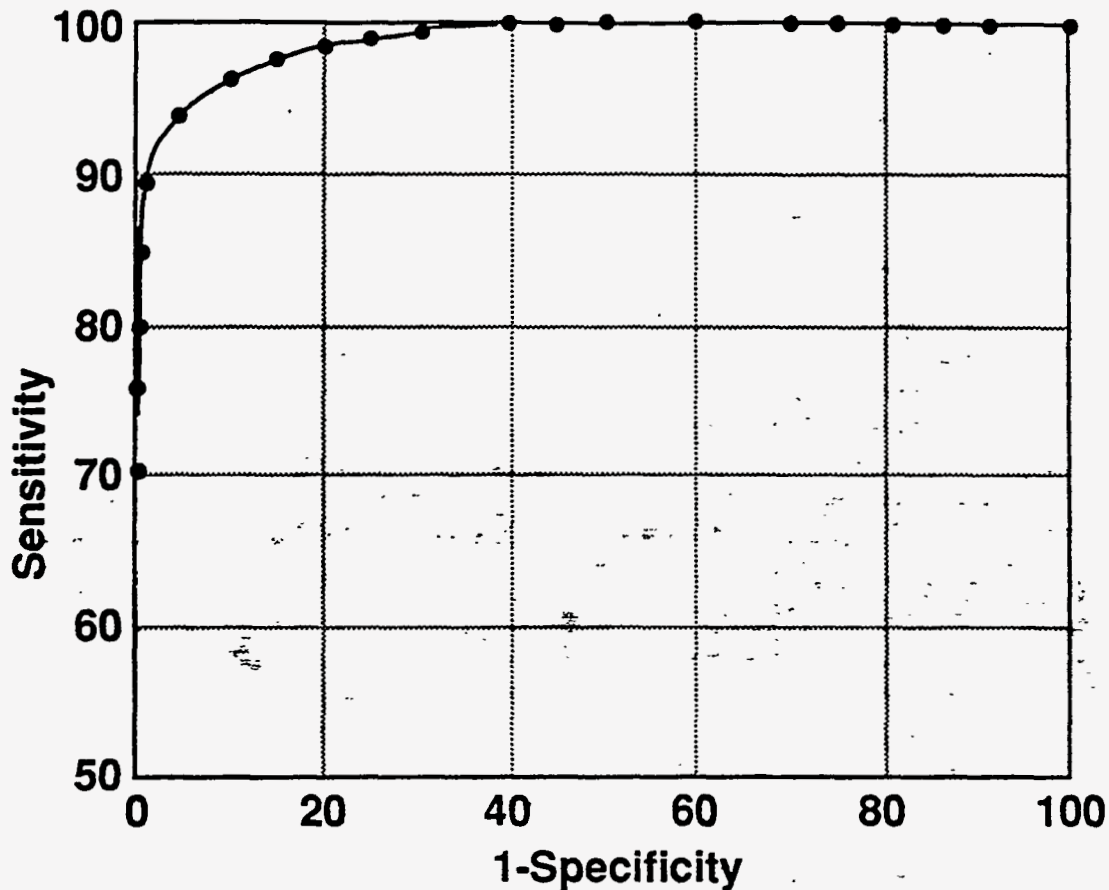


Figure 5. ROC display of the results of a Fisher Linear Discriminant classifier for features 7 and 12, (energy in the 20–21 and 21–22 KHz region.)

The previous study demonstrated the feasibility of classifying heart valve condition from acoustic signatures, we expanded the feature space and extended the data set to include test data. A neural network classifier was developed on the previously mentioned nine, confirmed data sets. Each of 21 additional unconfirmed heart valve data was processed in a hold one out procedure to test the performance of the neural network classifier. These unconfirmed data sets had been classified SLS maximum negative by x-ray analysis. A narrower frequency band (93 Hz) was considered and the optimal features were selected. Our results for these heart valves which are considered intact, are shown in Figure 7, (page 13). There is a correlation between the performance of the neural network classifier and the quality of the acoustic data. Noisier acoustic signals from valves 5 - 15 did not produce the expected classification results. We do not know if the mis-call is caused by the noise in the data or if the valves were truly SLS.

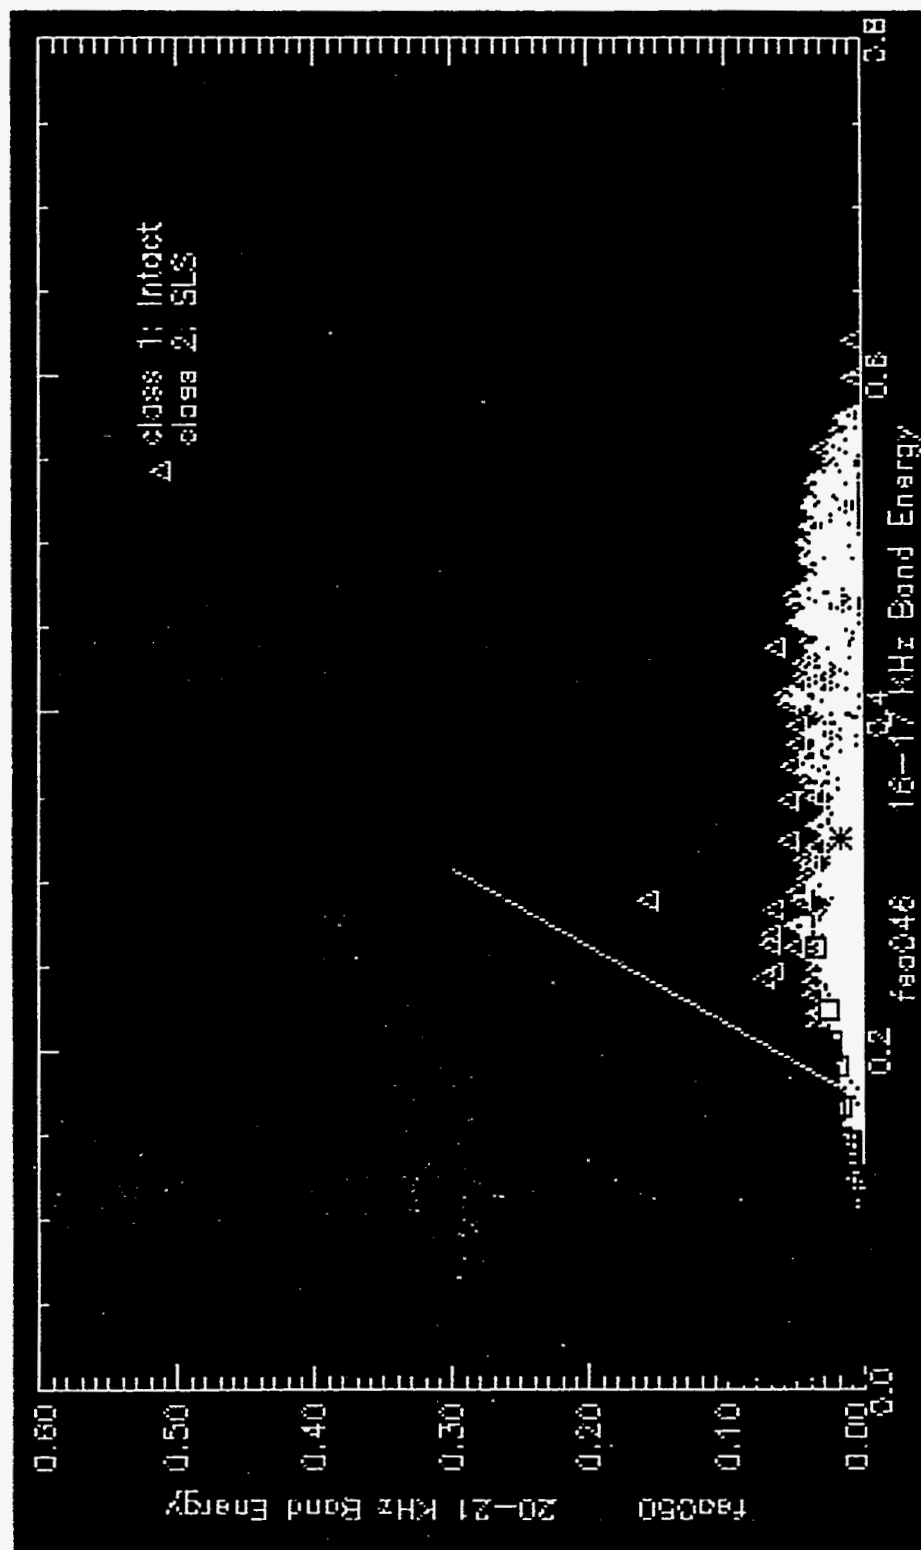


Figure 6. Cluster diagram displaying natural grouping of acoustic data for features 7 and 11.

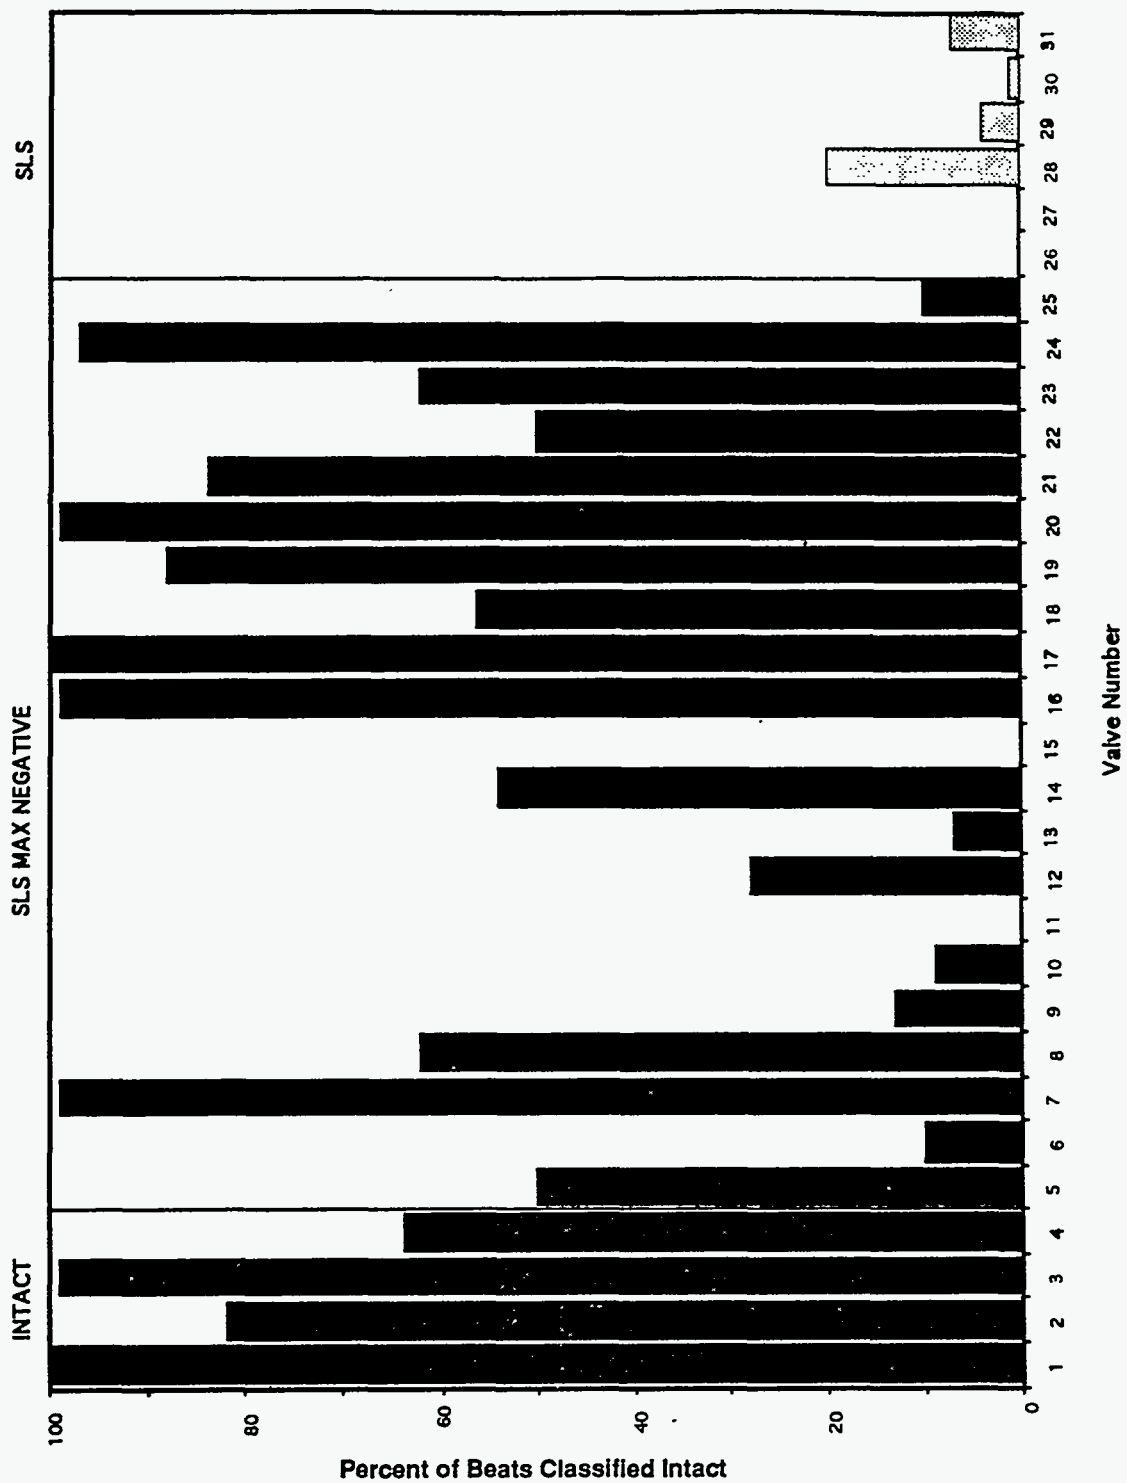


Figure 7. Neural network results for 93 Hz wide frequency features (phase II). For example, a threshold set at 30% would miss-call 7 of the 21 SLS maximum negative by x-ray analysis values.

These previously described studies bracketed the extremes in feature selection (phase I was 1 KHz wide frequency bands and phase II was 93 Hz wide frequency bands). The more robust features would fall in between these band widths. The third study we conducted was to determine the optimal features for classifying clinical data. The spectral features and the reflection coefficients were both extracted because previous research had indicated that these two features separated the intact and SLS valves relatively well. The spectral features consist of power centered at various frequencies in spectral bins of varying widths. Because the sampling frequency was 48 KHz and the number of points in each opening was 512, the smallest spectral bin width was 93.75 Hz ($48000 / 512$). Along with this bin width, bin widths of 187.50 Hz, 281.25 Hz, 375.00 Hz, 468.75 Hz, 562.50 Hz, 656.25 Hz, 750.00 Hz, 843.75 Hz, 937.50 Hz, 1031.25 Hz were also included. These are all multiples of the smallest bin (93.75 Hz). The reflection coefficients were generated using Burg's algorithm with a 20th order model. The data used for the training phase consists of features extracted from explanted valves. For the normal heart condition cases, there were four intact valves and six SLS valves. For the case of all heart conditions, there were four intact valves and seven SLS valves. In an attempt to avoid bias, an equal number of openings (or feature vectors) were chosen for each valve. The number decided upon was 100 openings per valve. Additionally, since our training set was very small, we decided to limit ourselves to only two features in the feature selection process.

The feature selection process determines which subset of features provide the greatest separation between the intact and the SLS valve classes. This is accomplished by computing a probabilistic distance measure between the classes. The feature subset which results in the greatest distance measure will tend to provide the greatest separation between the classes. The algorithm implemented to accomplish this is the sequential forward selection algorithm with the Mahalanobis distance. Because the sequential forward selection algorithm is sub-optimal, it does not always select the features subsets which provide the largest distance measure. For this reason, the sequential forward selection algorithm was used as a first step in the feature selection process. The second step was to analyze one-dimensional and two-dimensional feature space distributions for features selected by the feature selection algorithm in an attempt to select features which provide the greatest separation between the intact and the SLS valves. The two features selected were the spectral power in a band whose width is 500 Hz centered at 22.6 KHz and the spectral power in a band whose width is 500 Hz centered at 16.6 KHz.

As an example of the feature space distributions, the one-dimensional probability density functions for the first case is displayed in Figure 8. The red distributions represent the SLS densities, and the green distributions represent the intact densities. The top graph display the density functions

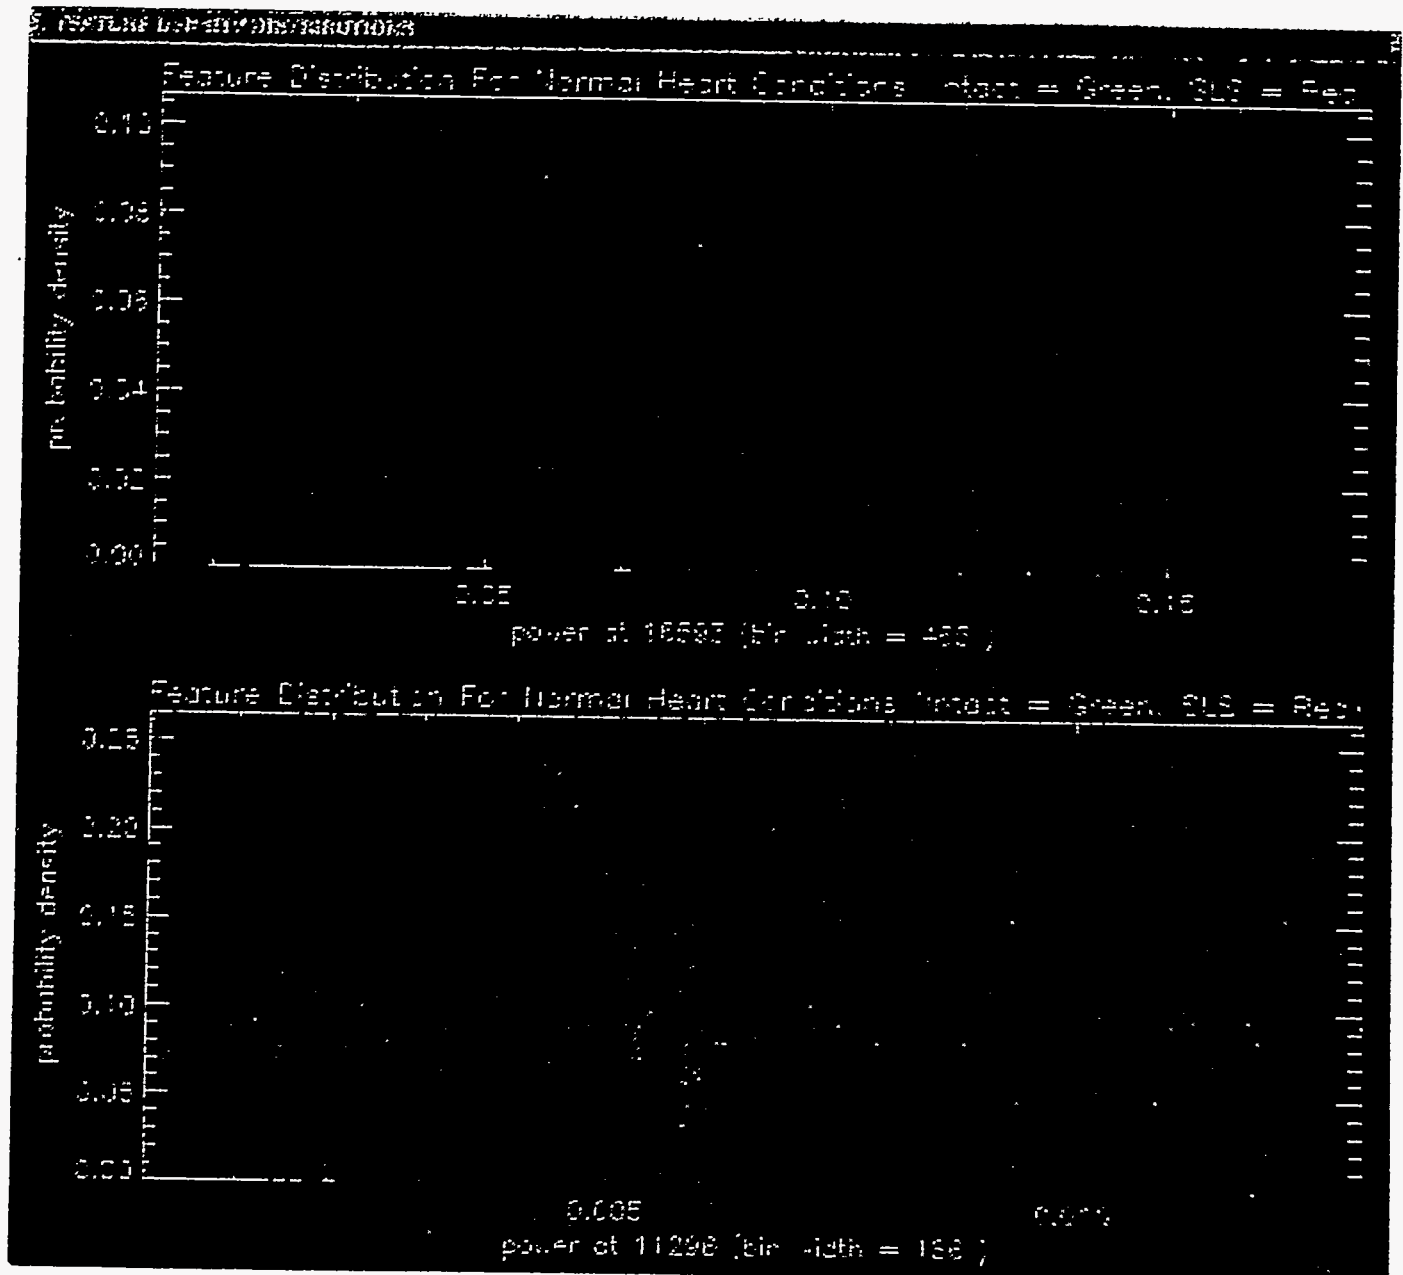


Figure 8. Probability density distribution for the two of the best features for this phase III study. This is part of the protocol for selecting feature for classification.

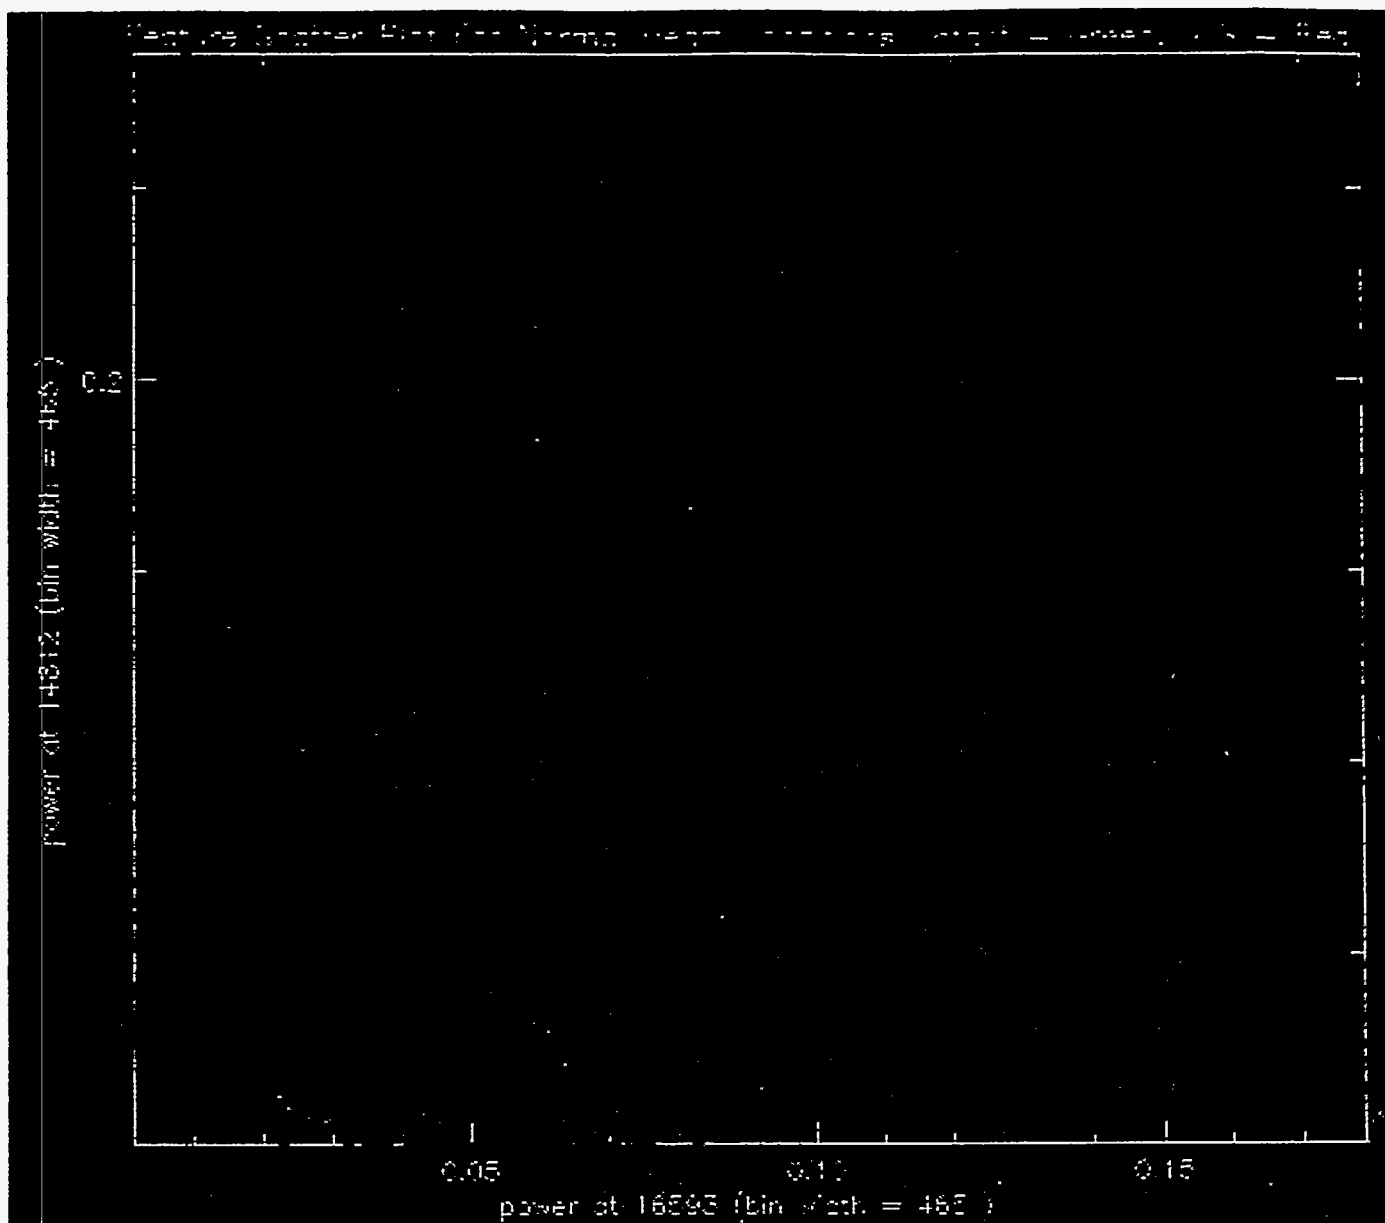


Figure 9. Cluster diagram for two optimized features for the phase III study. The natural clustering indicates classification potential for these features.

for the power centered at 16.6 KHz within a bandwidth of 500 KHz. The bottom graph display the density functions for the power centered at 11.3 KHz within a bandwidth of 200 KHz. From these graphs, class separation is clearly apparent with the upper plot showing better separation. The two-dimension feature space distributions for these features are shown in Figure 9. Again, the red represents the SLS features, and the green represents the intact features. For this figure, the separation between the classes is again evident.

After the feature selection has been performed, a probabilistic neural network is tuned using the "hold-one-out" procedure (Crawford, 1993, Buhl, 1993). This procedure trains a neural network with the chosen features for all training valves, except for one which has its feature vectors held out. Next, the feature vectors for the "held-out" valve are classified. This procedure is repeated for all valves and a sensitivity and a specificity are computed. If the results are unsatisfactory, this process is repeated using different neural network parameters (Buhl, 1993). This continues until the neural network has been tuned. These tuned neural networks then classify the "blinded" data set. The results of this process are shown in Figure 10, for valve numbers 1-4 and 26-31.

When the training phase produces acceptable results, the classification procedure is tested on a separate data set. Features determined in the training phase along with the "tuned" neural networks classify all the feature vectors extracted from the valves of the unknown condition. After this has been done, a percentage is computed which represents the percent of feature vectors classified as SLS. If the percentage of SLS vectors for a single heart valve is greater than 50 percent then the valve is called SLS. The higher the percentage of SLS calls for a particular valve, the greater the confidence in the prediction. We tested our classification algorithm on the "SLS max neg" data set which had not been confirmed by explant. These results are shown in Figure 10 for valves 5-25.

Conclusions

We have demonstrated that the acoustic signatures of the functioning heart valve provides information as to the condition of the outlet strut. Our research is based on the opening sounds. We have developed advanced signal processing techniques to extract the opening 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 clinical data.

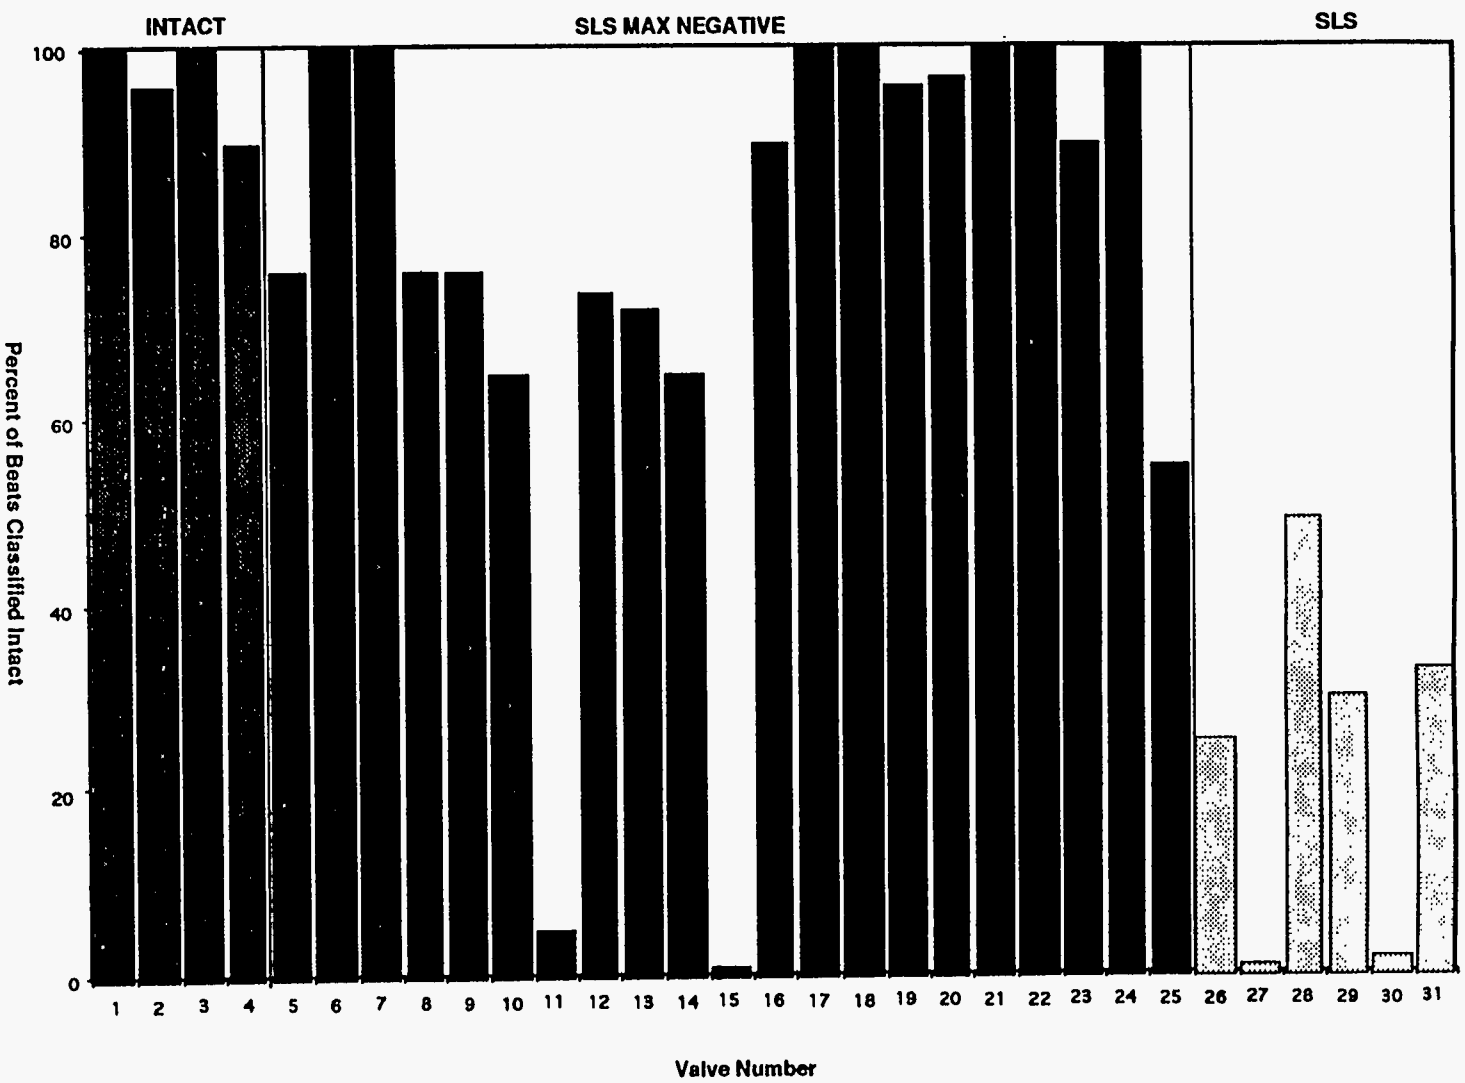


Figure 10. Neural network classification results for optimized features (16.6 KHz and 11.3 KHz @ 186 Hz [Phase III]). As an example, a threshold of 50% would miss-call two of the 21 SLS maximum negative by x-ray analysis valves with this improved classifier.

Acknowledgments

The authors would like to thank Jim Candy and Holger Jones for developing the signal processing algorithms and supplying the processed data for these classification efforts. We would especially like to thank David Wieting, Ph.D. and Rebecca Inderbitzen, M.S. of the Shiley Heart Valve Research Center for supporting and guiding our research.

This work was performed under the auspices of the Department of Energy by the Lawrence Livermore National Laboratory under contract W-7405-Eng-48.

References

Azevedo, S., Candy, J., Lager, D., On-line failure detection of vibrating structures. ASME Confr. Mech. Vibr. Noise (1981)

Ballard Dana H, Brown Christopher M. Computer Vision. Englewood, NJ: Prentice-Hall; 1982.

Buhl, Michael R., et al, "Detection of 'Single-Leg Separated' Heart Valves Using Statistical Pattern Recognition With the Nearest Neighbor Classifier", UCRL-ID-114802, July 16, 1993.

Candy, J. , Rozsa, R., Safeguards for a plutonium concentrator—an applied estimation approach. Automatica. (1980) 16 (6), 615-627.

Candy, J. Signal Processing: The Modern Approach. New York, N.Y.: McGraw Hill, (1988).

Candy, J., Barnes, F., Heart valve processing: a feasibility study. LLNL Report. (1991) UCRL-ID-107630.

Candy, J., Jones, H. E., Processing of Prosthetic Heart Valve Sounds for Single Leg Separation Classification, to be published

Crawford, S. L., and Thomas, G. H., "In-Vivo Classification of the Bjork Shiley Convexo-Concave Heart Valve from Acoustic Signatures", UCRL-ID-114819, August 1, 1993.

del Grande NK, Clark GA, Durbin PF, Fields DJ, Hernandez JE, Sherwood RJ. Buried Object Remote Detection Technology for Law Enforcement. SPIE Orlando '91 Symposium, Orlando, Florida, April 1-5, 1991.

Duda RO, Hart PE. Pattern Classification and Scene Analysis. New York, NY: Wiley; 1973.

Haykin, S. Adaptive Filter Theory. Englewood, N.J.: Prentice-Hall, (1986).

Jain AK. Fundamentals of Digital Image Processing. Englewood, NJ: Prentice-Hall; 1989.

Kohonen T. Self-Organization and Associative Memory. New York, NY: Springer-Verlag; 1989.

Ljung, L. System Identification: Theory for the User. Englewood, N.J.: Prentice-Hall, (1987).

Marr D. Vision. New York, NY. W. H. Freeman and Co; 1982.

Mullenhoff, Carmen, "Signal Processing of Shiley Heart Valve Data for Fracture Detection", UCRL-ID-113760, April 1, 1993.

Oppenheim, A., Shafer, R. Discrete-Time Signal Processing. Englewood, N.J.: Prentice-Hall, (1989).

Orphanidis, S. Optimum Signal Processing: An Introduction. New York, N.Y.: Macmillan, (1985).

Pratt WK. Digital Image Processing. New York, NY: Wiley; 1978.

Rosenfeld A. Computer Vision: Basic Principles. Proceedings of the IEEE.

1988; 76:863-8.

Rumelhart, D. E., and McClelland, J.L., "Parallel Distributed Processing: Explorations in the Microstructure of Cognition" Vol 1: Foundations, MIT Press, 1986.

Specht, D. F., "Probabilistic Neural Networks," Vol 3, Neural Networks, 1990.

Soderstrom, T., Stoica, P. System Identification. Englewood, N.J.: Prentice-Hall, (1989).

Weschler H. Computational Vision. San Diego, CA: Academic Press; 1990.