

1 of 1

1/2
8/16/93 JES (U)

UCRL-ID-114802

**Detection Of "Single-Leg Separated" Heart Valves
Using Statistical Pattern Recognition
With The Nearest Neighbor Classifier**

**Michael R. Buhl
Gregory A. Clark
James V. Candy
Graham H. Thomas**

July 16, 1993

 Lawrence
Livermore
National
Laboratory

This is an informal report intended primarily for internal or limited external distribution. The opinions and conclusions stated are those of the author and may or may not be those of the Laboratory.

Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

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.

**This report has been reproduced
directly from the best available copy.**

**Available to DOE and DOE contractors from the
Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (615) 576-8401, FTS 626-8401**

**Available to the public from the
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.,
Springfield, VA 22161**

ABSTRACT

Detection Of "Single-Leg Separated" Heart Valves Using Statistical Pattern Recognition With The Nearest Neighbor Classifier

The goal of this work was to detect "single-leg separated" Björk-Shiley Convexo-Concave heart valves which had been implanted in sheep. A "single-leg separated" heart valve contains a fracture in the outlet strut resulting in an increased risk of mechanical failure. The approach presented in this report detects such fractures by applying statistical pattern recognition with the nearest neighbor classifier to the acoustic signatures of the valve opening. This approach is discussed and results of applying it to real data are given.

MASTER

REPRODUCTION OR TRANSFER OF THIS DOCUMENT IS UNLAWFUL

ACKNOWLEDGMENTS

I thank Jim Candy for his expertise in signal processing and program management, Greg Clark for his expertise in pattern recognition and program management, Graham Thomas for program management, and Carmen Mullenhoff, Tony DeGroot, and Bob Searfus for the opening extraction techniques. Lastly, I would like to thank Becky Inderbitzen of Shiley for providing us with the heartbeat data.

TABLE OF CONTENTS

| | |
|--|-----|
| ABSTRACT | i |
| ACKNOWLEDGMENTS | ii |
| TABLE OF CONTENTS | iii |
| | |
| Section 1 Introduction | 1 |
| | |
| Section 2 Statistical Pattern Recognition Process | 2 |
| 2.1. Pre-Processing | 5 |
| 2.1.1. Event Detection | 5 |
| 2.1.2. Bad Event Rejection | 7 |
| 2.1.3. Event Extraction..... | 8 |
| 2.1.4. Filter and Resample | 8 |
| 2.1.5. Opening Response Extraction | 9 |
| 2.1.6. Opening Response Normalization | 9 |
| 2.2. Feature Extraction | 9 |
| 2.3. Classification | 10 |
| | |
| Section 3 Results | 11 |
| 3.1. Low Bandwidth (1 kHz to 5 kHz) Opening Responses | 14 |
| 3.2. Middle Bandwidth (3 kHz to 7 kHz) Opening Responses..... | 21 |
| 3.3. High Bandwidth (7 kHz to 24 kHz) Opening Responses | 27 |
| | |
| Section 4 Conclusions | 32 |
| | |
| References | 35 |

| | |
|--|-----------|
| Appendix A Feature Vector Description..... | 36 |
| A.1. Autoregressive (AR) Coefficients | 37 |
| A.2. Autoregressive (AR) Spectral Density | 38 |
| A.3. Autoregressive, Moving-Average (ARMA) Coefficients | 40 |
| A.4. Autoregressive, Moving-Average (ARMA) Spectral Density..... | 41 |
| A.5. Reflection Coefficients | 43 |
| A.6. First-Order Statistics Which Describe the Shape of the Power Spectral Densitites | 44 |

1. Introduction

Artificial heart valves have long been a godsend for many patients. With an artificial valve replacing a diseased or damaged one, a patient may live a long and productive life. Unfortunately, mechanical failure of such valves can be fatal if not detected early [1, 2]. One such valve which has demonstrated an increased risk of failure is the Björk-Shiley Convexo-Concave heart valve. Between 1979 and November 1986, this valve had been implanted in about 86,000 patients worldwide. Of these patients, nearly 300 have died as a result of mechanical failure [1]. Concern about these failures has lead to the research described here and elsewhere. Our overall goal is to develop techniques for non-invasive detection of single-leg separated heart valves. While other organizations are using a variety of other techniques, including the use of x-ray images, our work is focused on analyzing acoustic temporal waveform signatures.

A diagram illustrating the operation of the Björk-Shiley Convexo-Concave heart valve is shown in Fig. 1.1.

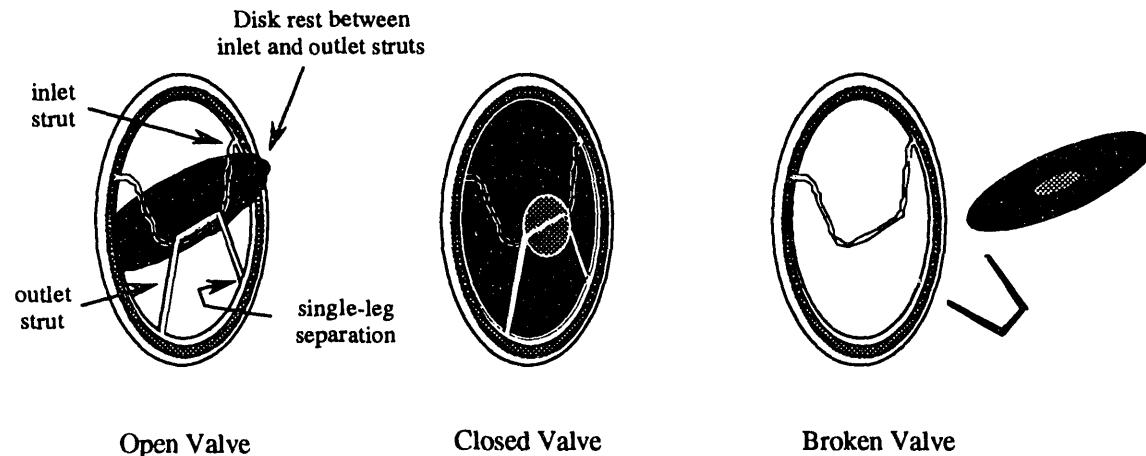


Figure 1.1. Björk-Shiley Convexo-Concave heart valve

Mechanical failure of these valves typically begins with a fracture, or a "single-leg separation" (SLS), which develops at the base of the outlet strut. After this condition

develops, there is an increased risk of complete detachment of the strut resulting in valve failure. Unfortunately, it is believed that the probability of death during replacement surgery is greater than the probability of mechanical failure. Therefore, patients with these valves have been advised not to have them replaced unless there is reason to believe their valve is "single-leg separated". If techniques are developed which can detect these fractures, suspect valves can be replaced and patients' lives ultimately saved.

The goal of this work was to develop a statistical pattern recognition technique which could successfully detect "single-leg separated" Björk-Shiley Convexo-Concave heart valves which had been implanted in sheep. Sheep were used for this study because their heart and chest cavities closely resemble those of humans. This approach classified valves based upon features which were generated from the acoustic responses of the heart valve opening. Only openings were analyzed because during an opening the disk directly impacts the top of the outlet strut. Therefore, the acoustic response of the opening should be closely coupled to the condition of this strut.

The remainder of this report is divided into two sections. Section two describes the statistical pattern recognition process and section three presents the results of applying this process to real heartbeat data.

2. Statistical Pattern Recognition Process

The process involved in the heart valve classification is shown in Fig. 2.1. The input into this process consists of heartbeats for the sheep whose valve condition is to be determined (test sheep), heartbeats for sheep with known intact valves, and heartbeats for sheep with known SLS valves. The output is a determination as to the condition of the valve implanted in the test sheep. The reader is referred to [3] for additional information about the measurements and signal pre-processing methods. Fig. 2.1 shows this process applied to three bandwidths of data as will be described below.

The first step performed on all the data was pre-processing. The goal of this step was to extract a manageable number of filtered opening responses which had relatively high signal-to-noise ratios. For this study, one-hundred responses were extracted for each of three frequency bands. Three bands were analyzed because the power spectrum of the responses was not uniformly distributed. Therefore, in order to study important parts of the spectrum effectively, the openings were filtered into bands, and each band was analyzed separately. In Fig. 2.1, the opening responses for each band are represented by an arrow

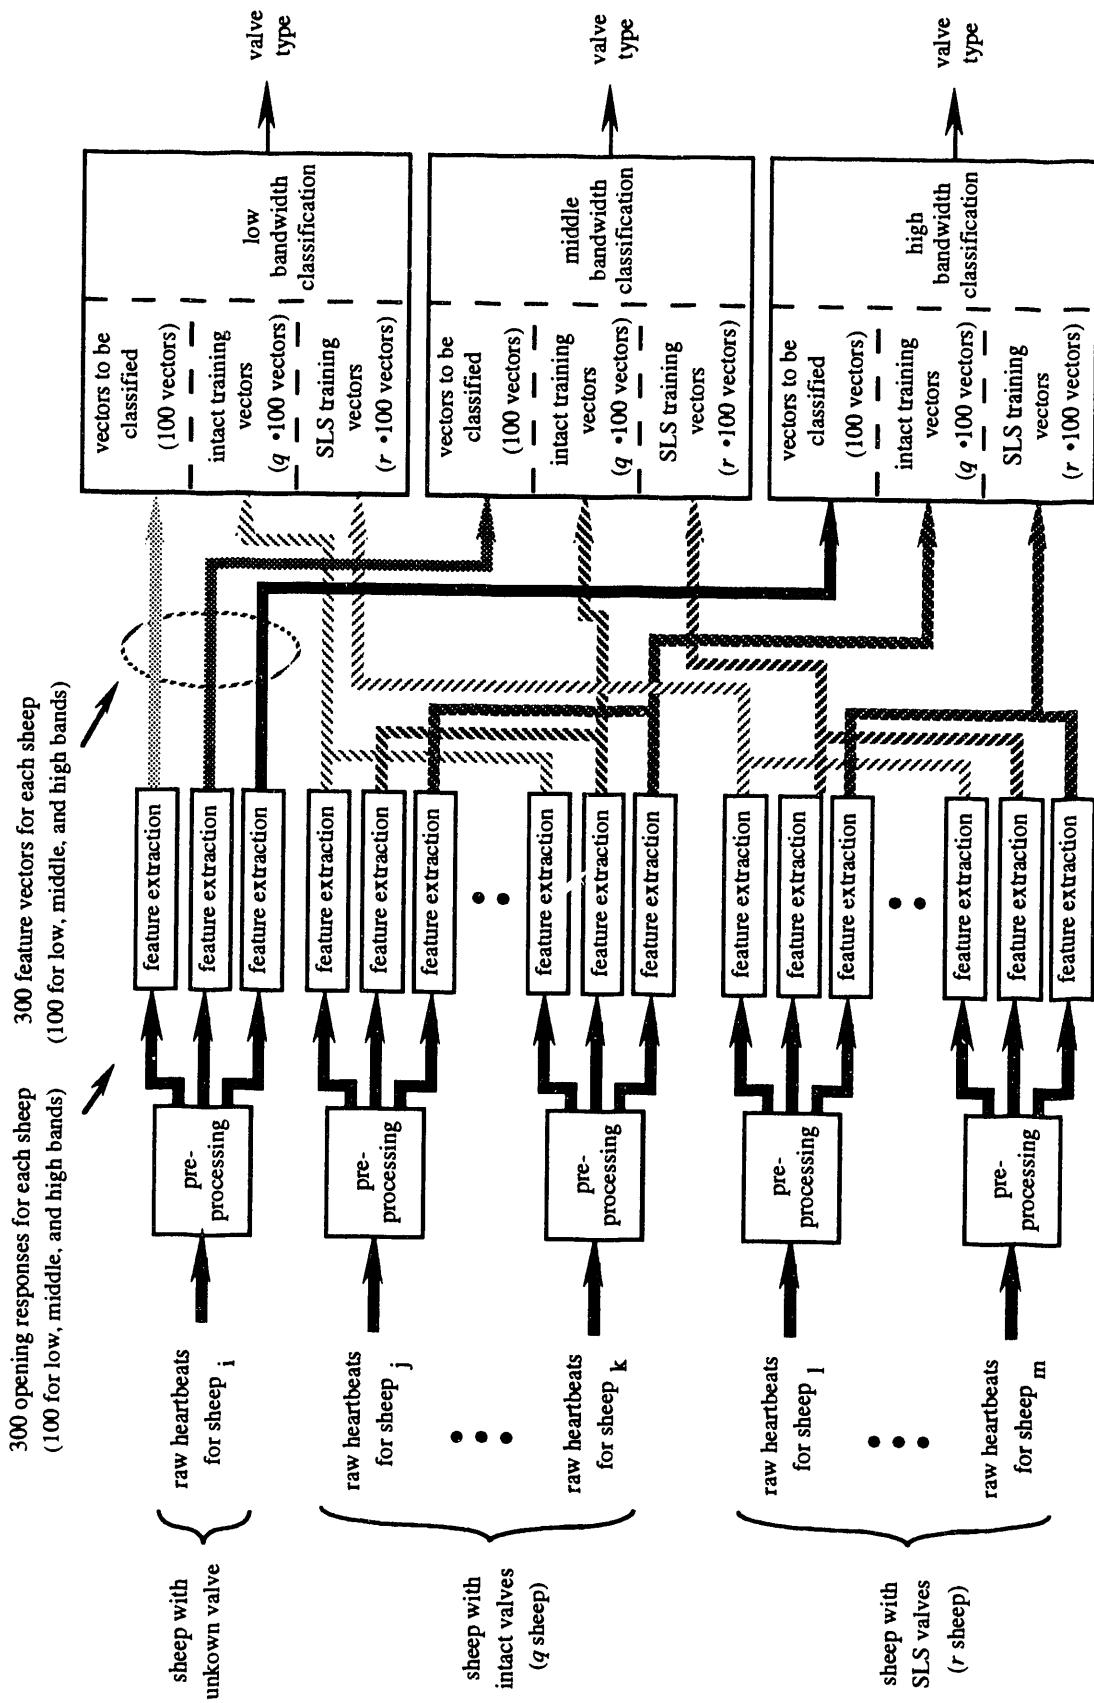


Figure 2.1. Block diagram of pattern recognition process for three bands of data

out of the appropriate pre-processing block. After pre-processing was completed, feature extraction was performed. The goal of this step was to generate a vector of features for each opening response. These features were used to represent various characteristics of these responses. Each feature extraction box, shown in Fig. 2.1, computed one-hundred feature vectors from one-hundred opening responses. After feature extraction had been performed for all sheep, the last step was to make a decision as to the condition of the valve which was implanted in the test sheep. This decision was based upon the comparison of the one-hundred feature vectors which had been generated for this sheep with the feature vectors which had been generated for the training sheep (sheep with known valves implanted).

Each of the three main operations presented above and shown in Fig. 2.1, pre-processing, feature extraction, and classification are described in more detail below.

2.1. Pre-Processing

The raw data files consisted of heartbeats, which in turn consisted of acoustic signatures of the valve closing and opening. The goal of this step was to process these files to create new files which contained a manageable number of filtered opening responses with relatively high signal-to-noise ratios. The steps involved in this process are shown in Fig. 2.2. Each of these steps is summarized in sections 2.1.1 to 2.1.6. For a more detailed description refer to [3].

2.1.1. Event Detection

The purpose of this step was to detect the temporal location and the maximum magnitude of events which represented possible closings and openings of the valve. This was accomplished in a series of steps. First, the raw heartbeats were bandpass filtered.

This operation smoothed the data and suppressed noise so the beats were better conditioned for the detection process. Then, from the filtered raw data, a signal which consisted of the

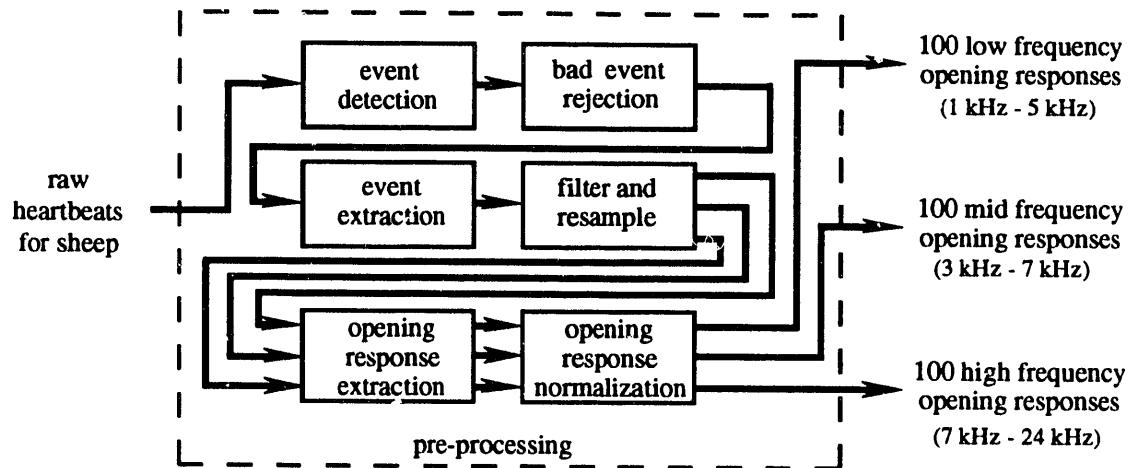


Figure 2.2. Pre-processing block diagram

ratio γ of the sliding short term average to the sliding long term average was computed. The short term average was computed over 50 samples, which represented 1.042 milliseconds (sampling period = 0.02083 milliseconds), and the long term average was computed over 500 samples, which represented 10.417 milliseconds. The resulting ratio γ was used to detect the location of events in the raw data. In order to do this, a threshold and a window size for γ were specified. The threshold specified the minimum ratio which was necessary for a sample point to be considered within an event, and the window size specified the maximum distance by which two sample points whose corresponding ratios exceeded the threshold could be separated and still be considered to be contained within the same event. Based upon experience with the data and engineering judgment, we chose to use a threshold of 3.0 and a window size of 2800 points. From these specifications, the beginning and ending locations, or sample points, of each event were determined. Next, the raw data files were used along with these locations to determine the maximum absolute signal value within each event. These values were eventually used to label the events as closings or openings. In order to do this, a threshold was automatically determined which

maximized the toggling between closings and openings. Using this threshold, all the events were labeled. Next, a beat time check which relied on the regularity of the beats was performed. The average time between similar events was calculated and a check was made for each event to determine whether it was within a twenty percent deviation from that average. When an out of range event was detected, it was marked as a "bad" event. Lastly, sequences consisting of at least five "good" (not "bad") events were written to a file. This file contained the beginning location of each event, the ending location of each event, the maximum absolute signal value within each event, and the type of each event (closing or opening). This information was written in the order in which the events occurred. The next step, bad event rejection, used this information to reject events which do not meet our criteria for useful signals, as defined next.

2.1.2. Bad Event Rejection

This step performed two operations on the file described above. First, the event ordering was checked to ensure a "open-close-open-close-open-etc" sequence. If more than one type of event was found in sequence, all events after the first were rejected. For example, if the sequence ordering was "close-open-open", the second opening was rejected. Next, the mean and standard deviation of the maximum absolute values for openings were computed. All close-open sequences were rejected if their maximum opening absolute value lay outside the interval defined by the mean $\pm n$ (standard deviation). If there were at least one-hundred beats (close-open sequence) detected in a file, $n = 0.5$, otherwise, $n = 1.0$. After this step, a file was written which contained the beginning location of each event, the ending location of each event, the maximum absolute signal value within each event, and the type of each event (closing or opening).

2.1.3. Event Extraction

This step used the file described above to extract the closings and openings from the raw data. It pasted these events into separate data files, one for closings and one for openings. Each event was centered in a window of 4096 points. Since the sampling period was 0.02083 milliseconds, this corresponded to a length of 85.33 milliseconds. This size was selected to ensure that the whole event was captured. For the results given in this report, only the openings were processed. Therefore, from this point on, only openings will be discussed. The next step was to filter and resample the openings.

2.1.4. Filter and Resample

When the power spectral estimates of the beats were analyzed, it was clear that the power of the spectrum was not distributed uniformly. The majority of the energy occurred at high frequencies. In order to effectively analyze important sections of the spectrum independently, the openings were filtered into three bands. The low band contained frequencies from 1 kHz to 5 kHz, the middle band contained frequencies from 3 kHz to 7 kHz, and the high band contained frequencies from 7 kHz to 24 kHz. After filtering, the low and middle band were down sampled to produce openings with fewer points. Hence, fewer computations were necessary in the feature extraction. For the low bandwidth, the data was down sampled by a factor of four, resulting in openings with 1024 points at a sampling period of 0.08333 milliseconds. For the middle bandwidth, the data was down sampled by a factor of two, resulting in openings with 2048 points at a sampling period of 0.04167 milliseconds. Lastly, for the high bandwidth resampling was not performed. Therefore, the openings were still 4096 points at a sampling period of 0.02083 milliseconds.

2.1.5. Opening Response Extraction

Each opening often consisted of multiple responses, which were the result of the blood flow past the disk causing the disk to hit the outlet strut multiple times. Moreover, the number of responses per opening was not consistent from beat to beat. This variable was assumed to be a result of the sheep's biological state, such as the sheep's blood pressure or pulse rate, and not the condition of the heart valve. Since each of these responses was caused by the same occurrence (flow of blood), each opening could be assumed to be a series of impulse responses of the disk hitting the strut. Therefore, for classification purposes, only one response needed to be analyzed at a time. The goal of this step was to inspect all the openings for each sheep, and select one-hundred opening responses of relatively high signal-to-noise ratios for each bandwidth. Because each response consisted of only a fraction of the opening, fewer points were needed to represent them. After this extraction, the low bandwidth responses consisted of 256 points, the middle bandwidth responses consisted of 512 points, and the high bandwidth responses consisted of 1024 points.

2.1.6. Opening Response Normalization

Before features were extracted, all the responses were normalized to unit variance. This step was necessary because the features which were generated are dependent upon the energy contained in each response.

2.2. Feature Extraction

The goal of this step was to generate a vector of features for each normalized opening response. Feature vectors represent in compressed form various characteristics which can be used to distinguish opening responses caused by SLS valves from opening responses caused by intact valves. Vectors which were analyzed for this study consisted of

coefficients for parametric models of the responses, power spectral densities, and first-order features which described the shape of the power spectral densities. After the feature vectors were generated for all sheep (testing and training), classification was performed.

2.3. Classification

The goal of this step was to make a decision as to the condition of the valve for the sheep which was being tested. This decision was based upon the comparison of the one-hundred feature vectors which had been generated for this sheep with the feature vectors which had been generated for the training sheep (sheep with known valves implanted). This classification was accomplished in two steps, as shown in Fig. 2.3.

First, each feature vector was classified as intact or SLS. A variety of classifiers could have been implemented. Some commonly used classifiers are the nearest neighbor classifier and neural networks. The nearest neighbor classifier was implemented because it is relatively simple, and there is no need to adjust parameters as may be the case for other classifiers including neural networks. This method computes the Euclidean distances among the feature vector which is to be classified and the training feature vectors. The feature vector is assumed to belong to the same class as the vector which produced the smallest distance [4].

The probabilistic neural network (PNN) is generally preferred because it is Bayes optimal and generally outperforms the nearest neighbor algorithm which is not [6]. The disadvantage of the PNN is that it requires choosing an optimal or acceptable value of a tuning parameter [6]. This has two important implications. First, a costly set of iterations must be performed with the data to choose (find) the optimal tuning parameter. Second, since we have a limited data set which is not large enough to give statistically significant results, we use the "hold one out method" of training and testing the classifier (see Section 3 of this report). This requires that we train the PNN M times for a given data set, where

M is the number of data samples, and each training requires iterations to choose the tuning parameter. We decided that this procedure is too costly when the "hold one out" method is required, so we avoided the cost by simply using the nearest neighbor algorithm. The nearest neighbor algorithm is suboptimal, but its simplicity offsets this disadvantage and the algorithm gives us good, useful results.

Once all one-hundred feature vectors had been classified for one of the bandwidths of opening responses, the condition of the valve was inferred based upon the percentage of feature vectors which had been classified as SLS. This specified percentage will be referred to as the "minimum percent SLS threshold". If the percentage of feature vectors classified as SLS is greater than or equal to this threshold, the valve is classified as SLS, otherwise the valve is classified as intact.

The next section describes results of applying the steps discussed in sections 2.1 to 2.3 to real data.

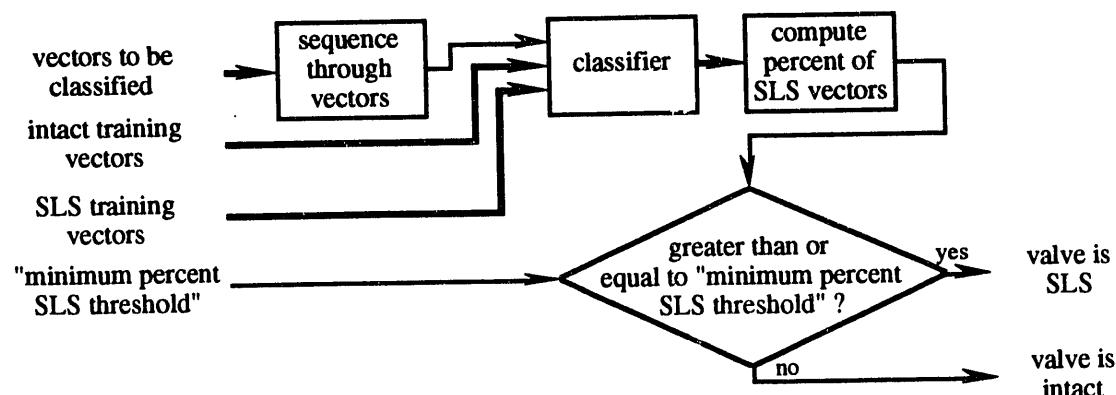


Figure 2.3. Classification block diagram

3 . Results

In order to analyze the effectiveness of this statistical pattern recognition approach in detecting SLS heart valves, a known data base had to be created. This data base, which was provided by Shiley, was created by implanting twelve sheep with intact valves and

seven sheep with SLS valves. Each of these sheep will be referred to by a number or in a few cases by letters preceding numbers. These sheep along with their valve conditions are given below.

Sheep with intact valves: 40, 74, 75, 95, 129, 209, 414, 4337, 4380, bf1, bf12, t520

Sheep with SLS valves: 68, 103, 105, 110, 411, 412, 422

Because the data base consisted of only nineteen sheep, the "hold one out" technique was used to test the effectiveness of this approach. This technique trains the classifier with data from all sheep except the sheep which was being tested. Therefore, when the testing sheep had an intact valve, the classifier was trained with 1100 feature vectors from the eleven other sheep with intact valves, and 700 feature vectors from the sheep with SLS valves. Likewise, when the testing sheep had a SLS valve, the classifier was trained with 1200 feature vectors from the sheep with intact valves, and 600 feature vectors from the six other sheep with SLS valves. Because three bandwidths of opening responses were tested, 57 classifiers (19 sheep • 3 bandwidths) were trained for each feature vector which was analyzed. Due to this large number, the nearest neighbor classifier was implemented as discussed in section 2.3. Other classifiers, such as neural networks, would require vast amounts of time to "tune" and train as discussed in the last section.

Several different feature vectors were analyzed for this study. These vectors are given below. For a more detailed description of these features refer to appendix A or [5].

- autoregressive (AR) coefficients
- AR power spectral densities
- autoregressive, moving-average (ARMA) coefficients
- ARMA power spectral densities
- reflection coefficients

- first-order statistics describing the shape of the ARMA power spectral densities
- first-order statistics describing the shape of the AR power spectral densities

Previous results indicated that the feature vectors which provided the best results were the AR power spectral estimates, the ARMA power spectral estimates, and the reflection coefficients. Therefore, these were the feature vectors which were extracted for the results presented below.

The AR model coefficients, which are used to construct the AR spectral density were computed using the Levinson-Durbin algorithm [5]. For the low bandwidth, fifteenth order models were used. For the middle bandwidth, twenty-fifth order models were used. Lastly, for the high bandwidth fiftieth order models were used. The orders for all these models were selected based on previous results.

The ARMA model coefficients, which are used to construct the ARMA spectral density, were computed using the recursive least-squares algorithm [5]. For the low bandwidth, sixth order AR and fourth order MA models were used. For the middle bandwidth, twelfth order AR and eighth order MA models were used. Lastly, for the high bandwidth, eighteenth order AR and twelfth order MA models were used. The orders for all these models were selected based on previous results.

The reflection coefficients were computed using Burg's algorithm [5]. For all bandwidths, fiftieth order models were computed. Again, this order was selected based on previous results.

One last point should be mentioned regarding the feature vectors. As a statistical rule of thumb for achieving statistically significant detection results, the maximum number of features per vector should be one-fifth of the minimum number of vectors (samples) available for a given class which are used to train the classifier. Since the minimum number of samples per class was 600 (when an SLS sheep is being tested), the maximum

number of features per vector should be 120 (600/5). Therefore, when the power spectral densities were extracted as features, they were represented as 120 frequency bins equally spaced between the bandwidth cut-off frequencies.

The next three sections present the results of applying the statistical pattern recognition process using the nearest neighbor classifier to the feature vectors for the low, middle, and high bandwidth responses.

3.1. Low Bandwidth (1 kHz to 5 kHz) Opening Responses

The results obtained with the low bandwidth opening responses are given in Table 3.1. This table indicates the percentage of feature vectors for each sheep which were classified as SLS. These percentages are based upon the one-hundred opening responses which were extracted per sheep. Instances where at least one-half of the feature vectors were classified as SLS are indicated in bold.

For the sheep with intact valves, the percentage of feature vectors classified as SLS are generally low. In fact, when the AR power spectral density was used as the feature vector, not one sheep had at least one-half of its vectors incorrectly classified as SLS. When the ARMA power spectral density was used as the feature vector, only one sheep (bf12) had at least one-half of its vectors incorrectly classified as SLS. Lastly, when the reflection coefficients are used as the feature vector, only two sheep (75 and 4337) had at least one-half of their vectors incorrectly classified as SLS.

For the sheep with SLS valves, the percentage of feature vectors classified as SLS are generally higher than for the sheep with intact valves. When both power spectral

| Sheep | Condition of Valve | Feature Vector | | |
|-------|--------------------|--|--|---|
| | | AR power spectral density (120 features) (%) | ARMA power spectral density (120 features) (%) | Reflection coefficients (50 features) (%) |
| 40 | intact | 16.00 | 15.00 | 22.00 |
| 74 | intact | 14.00 | 20.00 | 29.00 |
| 75 | intact | 47.00 | 42.00 | 68.00 |
| 95 | intact | 8.00 | 22.00 | 20.00 |
| 129 | intact | 12.00 | 14.00 | 19.00 |
| 209 | intact | 16.00 | 13.00 | 18.00 |
| 414 | intact | 5.00 | 14.00 | 15.00 |
| 4337 | intact | 33.00 | 25.00 | 68.00 |
| 4380 | intact | 6.00 | 5.00 | 38.00 |
| bf1 | intact | 0.00 | 1.00 | 0.00 |
| bf12 | intact | 32.00 | 51.00 | 23.00 |
| t520 | intact | 10.00 | 7.00 | 17.00 |
| 68 | SLS | 96.00 | 71.00 | 89.00 |
| 103 | SLS | 79.00 | 72.00 | 86.00 |
| 105 | SLS | 16.00 | 49.00 | 36.00 |
| 110 | SLS | 92.00 | 83.00 | 92.00 |
| 411 | SLS | 14.00 | 11.00 | 30.00 |
| 412 | SLS | 98.00 | 87.00 | 98.00 |
| 422 | SLS | 24.00 | 38.00 | 60.00 |

Table 3.1. Percentage of features vectors which were classified as SLS for low band opening responses

densities were used as feature vectors, four sheep (68, 103, 110, and 412) had at least one-half of their vectors correctly classified as SLS. When the reflection coefficients were used as the feature vectors, five sheep (68, 103, 110, 412, and 422) had at least one-half of their vectors correctly classified as SLS.

In order to classify the valves, one more step must be taken. This step is to specify a percentage which must be exceeded for a valve to be classified as SLS, as shown in Fig. 2.3. This threshold, which is referred to as the "minimum percent SLS threshold", provides a "knob" by which the "probability of detection" and the "probability of false alarm" can be varied. The "probability of detection" is defined as the percentage of SLS valves which were classified correctly as SLS, and the "probability of false alarm" is defined as the percentage of intact valves which were incorrectly classified as SLS. Therefore, if the "minimum percent SLS threshold" is very small, the "probability of detection" will be high, but the "probability of false alarm" will also be high. Likewise, if the "minimum percent SLS threshold" is very large, the "probability of false alarm" will be low, but the "probability of detection" will also be low. Ideally, it is desirable to have a high "probability of detection" and a small "probability of false alarm". The trade-offs which exist by varying this threshold are shown in Fig's. 3.1 to 3.3. Fig. 3.1 illustrates the probabilities which can be achieved by using the AR power spectral density as the feature vector, Fig. 3.2 illustrates the probabilities which can be achieved by using the ARMA power spectral density as the feature vector, and Fig. 3.3 illustrates the probabilities which can be achieved by using the reflection coefficients as the feature vector. The top graph in each of these figures displays the "probability of detection" as a function of the "minimum percent SLS threshold", and the bottom graph displays the "probability of false alarm" as a function of the "minimum percent SLS threshold".

Table 3.2 presents the best range of operating points from the graphs which result in "probabilities of detection" of at least 71.43% and "probabilities of false alarm" of at most

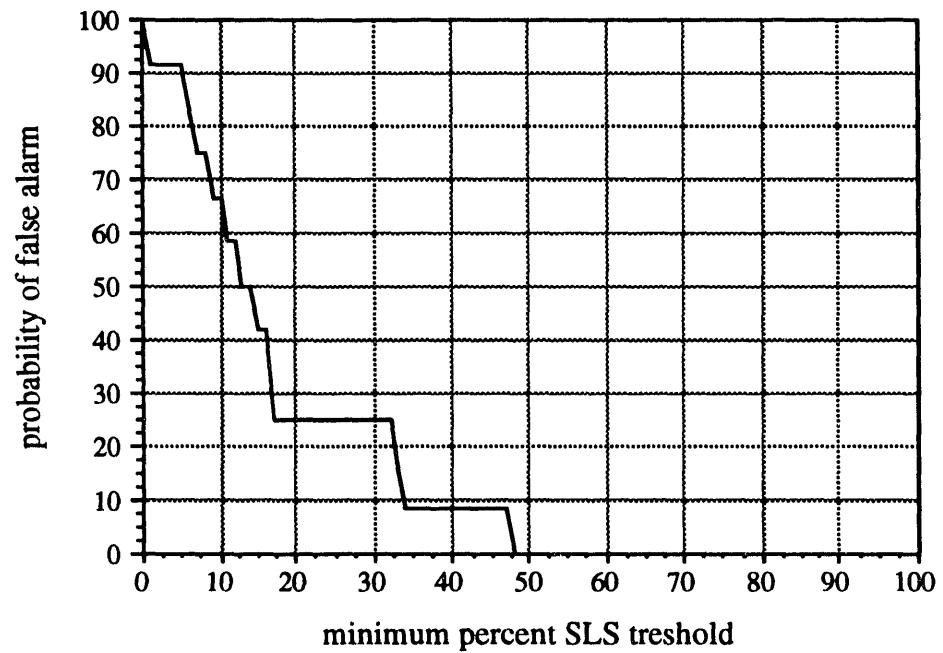
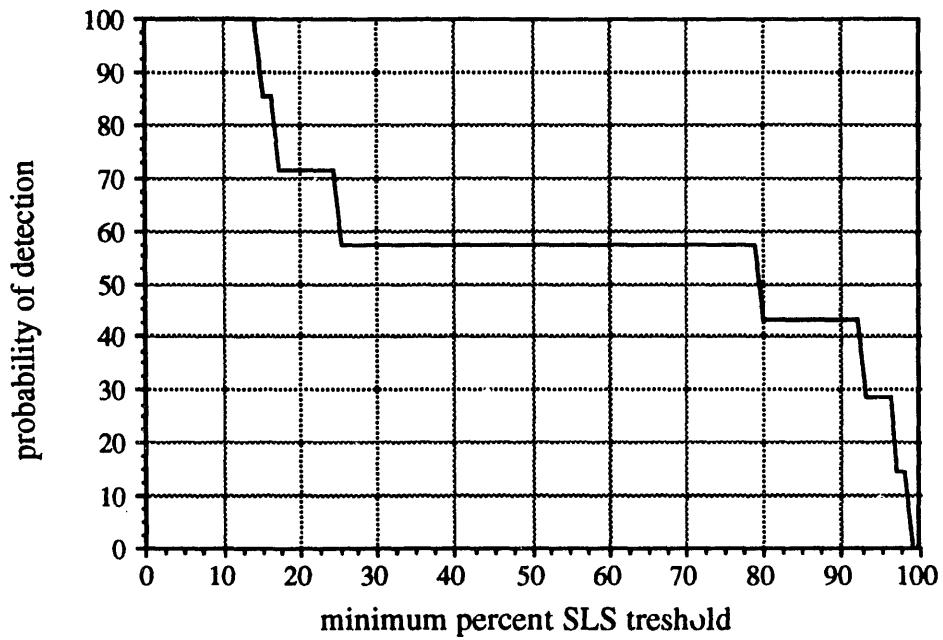


Figure 3.1. "Probability of detection" and "probability of false alarm" versus "minimum percent SLS threshold" for classification using AR power spectral density for low band openings

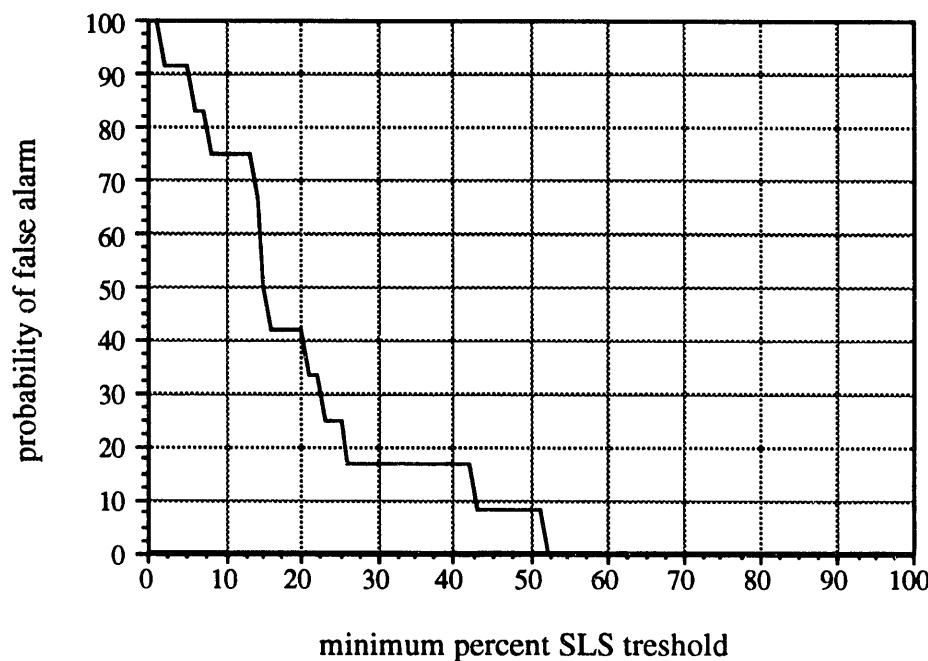
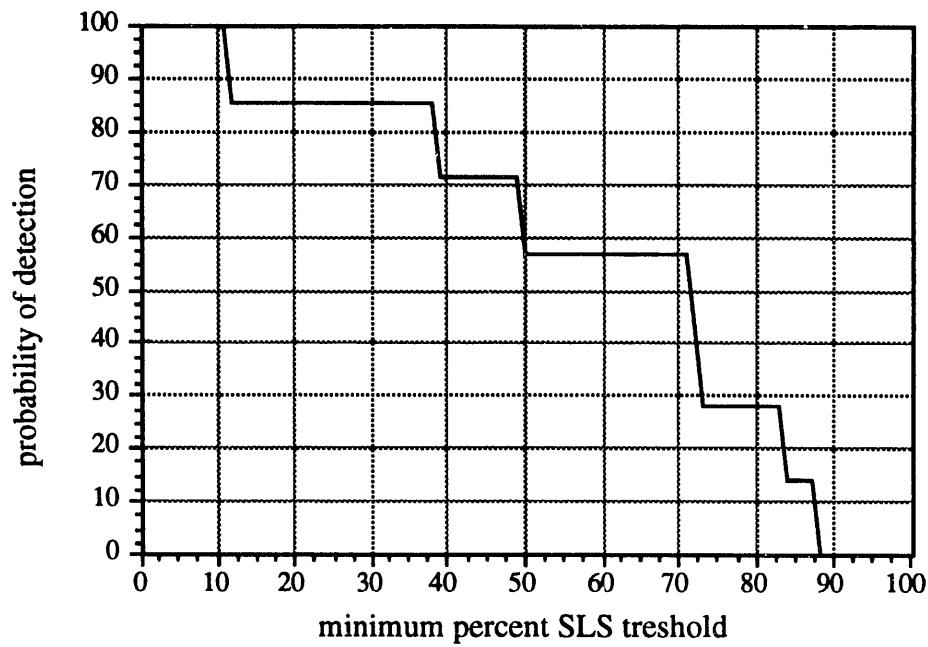


Figure 3.2. "Probability of detection" and "probability of false alarm" versus "minimum percent SLS threshold" for classification using ARMA power spectral density for low band openings

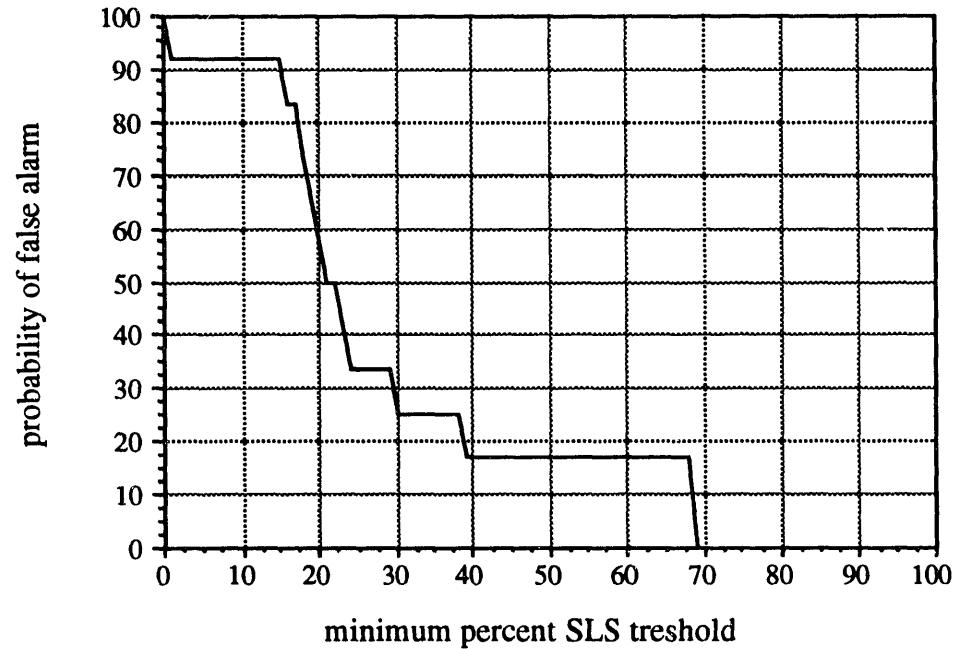
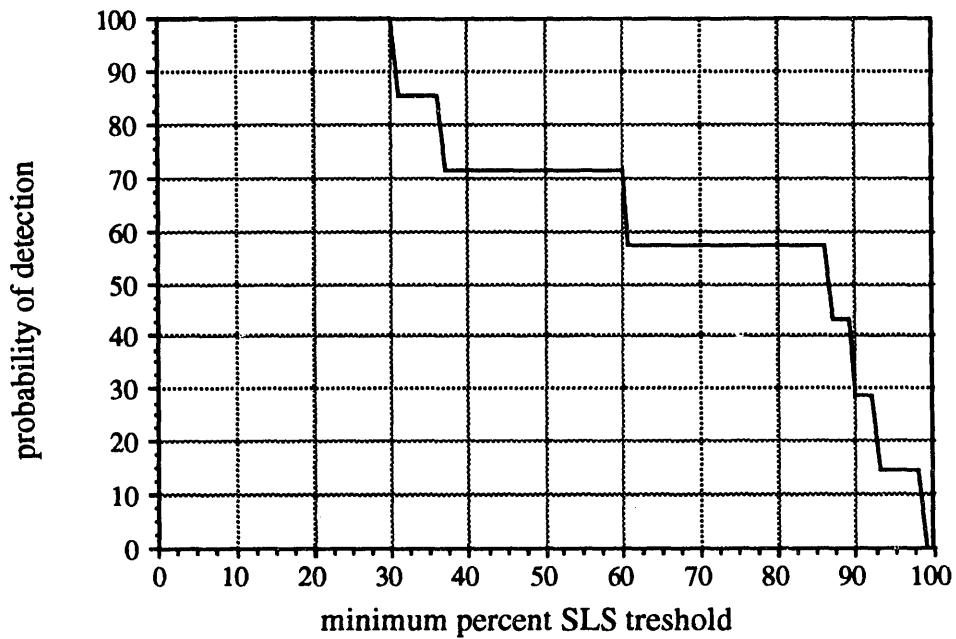


Figure 3.3. "Probability of detection" and "probability of false alarm" versus "minimum percent SLS threshold" for classification using reflection coefficients for low band openings

| Feature Vector | Minimum Percent SLS Threshold | Probability of Detection | Probability of False Alarm |
|-----------------------------|-------------------------------|--------------------------|----------------------------|
| AR power spectral density | 17.00 - 24.00 | 71.43 (5/7) | 25.00 (3/12) |
| ARMA power spectral density | 26.00 - 38.00 | 85.71 (6/7) | 16.67 (2/12) |
| ARMA power spectral density | 43.00 - 49.00 | 71.43 (5/7) | 8.33 (1/12) |
| reflection coefficients | 30.00 | 100.00 (7/7) | 25.00 (3/12) |
| reflection coefficients | 39.00 - 60.00 | 71.43 (5/7) | 16.67 (2/12) |

Table 3.2. "Probabilities of detection" and "probabilities of false alarm" which can be achieved from the low band opening responses

25.00%. From this table, it appears that the ARMA power spectral density and the reflection coefficients seem to provide higher "probabilities of detection" and lower "probabilities of false alarm" than does the AR power spectral density. The decision as to which feature vector and which "minimum percent SLS threshold" give the best results depends on the specifications which are desired. From the data which were used in this study, it was possible to achieve a "probability of detection" of 100% with a "probability of false alarm" of 25% by using the reflection coefficients with a threshold of 30%. If a lower false alarm rate is desired, this can be achieved by using the ARMA power spectral density with a threshold from 26% to 38%. For this case, the "probability of detection" is 85.71% and the "probability of false alarm" is 16.67%. If it is desirable to have an even lower false alarm rate, the ARMA power spectral density can be used with a threshold from 43% to 49%. For this case, the "probability of detection" is 71.43% and the "probability of false alarm" is 8.33%.

3.2. Middle Bandwidth (3 kHz to 7 kHz) Opening Responses

The results obtained with the middle bandwidth opening responses are given in Table 3.3. This table indicates the percentage of feature vectors for each sheep which were classified as SLS. These percentages are based upon the one-hundred opening responses which were extracted per sheep. Instances where at least one-half of the feature vectors were classified as SLS are indicated in bold.

For the sheep with intact valves, the percentage of feature vectors classified as SLS are again generally low. When the AR power spectral density was used as the feature vector, only two sheep (75 and 414) had at least one-half of their vectors incorrectly classified as SLS. When the ARMA power spectral density was used as the feature vector, again only two sheep (414 and bf12) had at least one-half of their vectors incorrectly classified as SLS. Lastly, when the reflection coefficients are used as the feature vector, three sheep (95, bf1, and t520) had at least one-half of their vectors incorrectly classified as SLS.

For the sheep with SLS valves, the percentage of feature vectors classified as SLS are generally higher than for the sheep with intact valves. When the AR power spectral density was used as the feature vector, five sheep (68, 103, 105, 110, and 422) had at least one-half of their vectors correctly classified as SLS. When the ARMA power spectral density was used as the feature vector, four sheep (103, 105, 110, and 422) had at least one-half of their vectors correctly classified as SLS. Lastly, when the reflection coefficients were used as the feature vector, six sheep (all except 68) had at least one-half of their vectors correctly classified as SLS.

Fig's. 3.4 to 3.6 illustrate the "probabilities of detection" and the "probabilities of false alarm" which are achievable by varying the "minimum percent SLS threshold". Fig. 3.4 illustrates achievable probabilities using the AR power spectral density as the feature

| Sheep | Condition of Valve | Feature Vector | | |
|-------|--------------------------|---|---|--|
| | | AR power spectral density (120 features) (%) | ARMA power spectral density (120 features) (%) | Reflection coefficients (50 features) (%) |
| 40 | intact | 21.00 | 18.00 | 18.00 |
| 74 | intact | 5.00 | 6.00 | 16.00 |
| 75 | intact | 52.00 | 34.00 | 33.00 |
| 95 | intact | 22.00 | 26.00 | 56.00 |
| 129 | intact | 35.00 | 37.00 | 14.00 |
| 209 | intact | 5.00 | 33.00 | 24.00 |
| 414 | intact | 53.00 | 54.00 | 39.00 |
| 4337 | intact | 12.00 | 18.00 | 2.00 |
| 4380 | intact | 36.00 | 36.00 | 24.00 |
| bf1 | intact | 18.00 | 41.00 | 60.00 |
| bf12 | intact | 48.00 | 60.00 | 17.00 |
| t520 | intact | 35.00 | 33.00 | 95.00 |
| 68 | SLS | 55.00 | 29.00 | 26.00 |
| 103 | SLS | 58.00 | 50.00 | 67.00 |
| 105 | SLS | 59.00 | 58.00 | 68.00 |
| 110 | SLS | 83.00 | 77.00 | 53.00 |
| 411 | SLS | 34.00 | 49.00 | 68.00 |
| 412 | SLS | 40.00 | 36.00 | 53.00 |
| 422 | SLS | 67.00 | 57.00 | 72.00 |

Table 3.3. Percentage of features vectors which were classified as SLS for middle band opening responses



Figure 3.4. "Probability of detection" and "probability of false alarm" versus "minimum percent SLS threshold" for classification using AR power spectral density for middle band openings

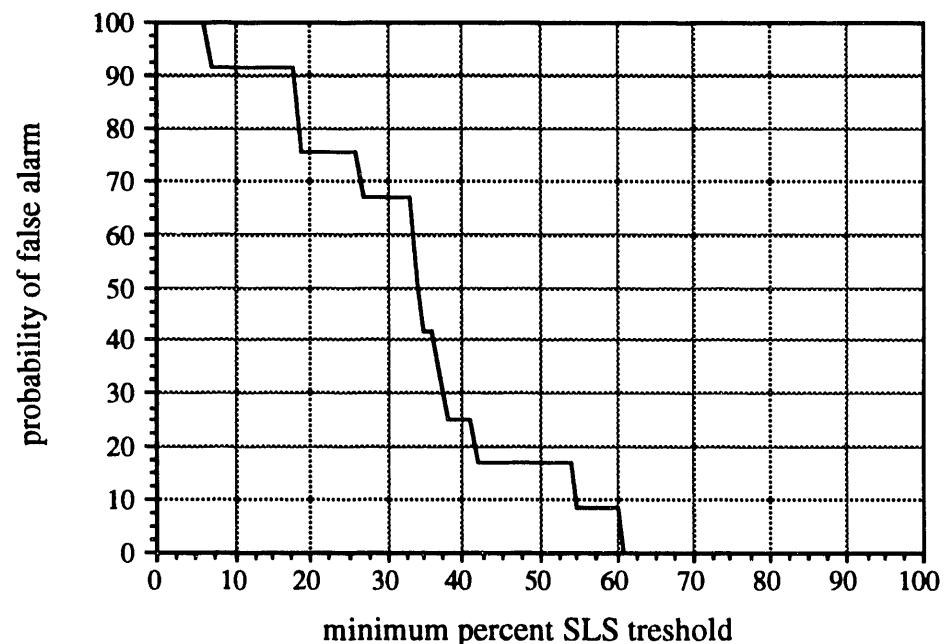
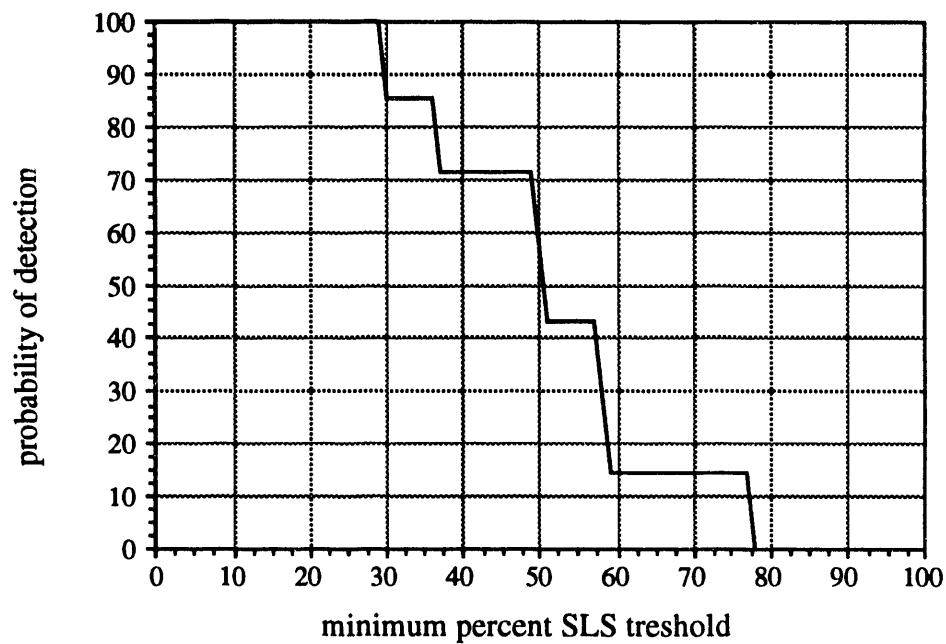


Figure 3.5. "Probability of detection" and "probability of false alarm" versus "minimum percent SLS threshold" for classification using ARMA power spectral density for middle band openings

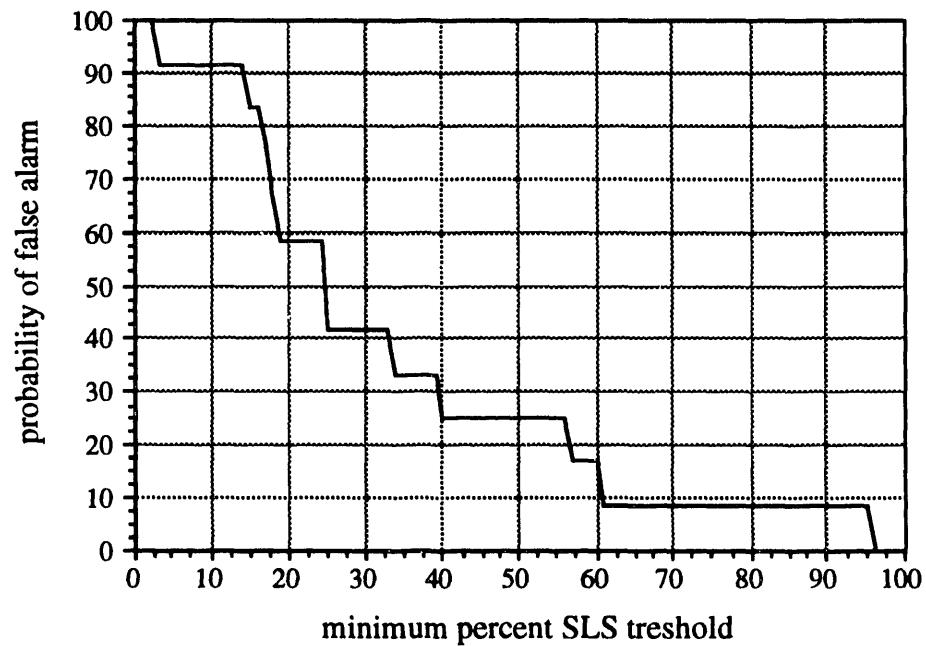
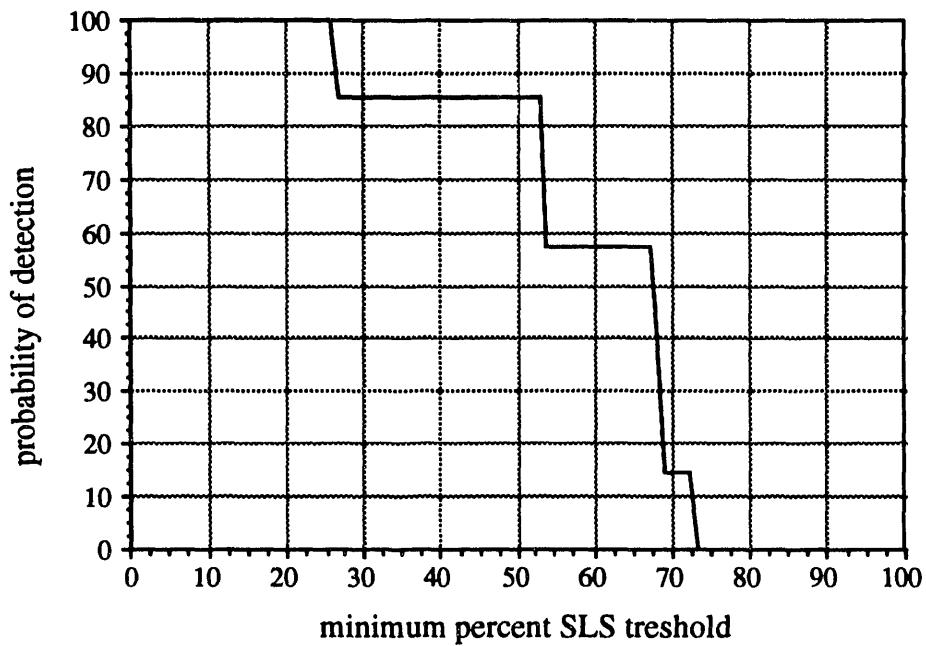


Figure 3.6. "Probability of detection" and "probability of false alarm" versus "minimum percent SLS threshold" for classification using reflection coefficients for middle band openings

vector, Fig. 3.5 illustrates achievable probabilities using the ARMA power spectral density as the feature vector, and Fig. 3.6 illustrates achievable probabilities using the reflection coefficients as the feature vector.

Table 3.4 present the best operating points from the graphs which result in "probabilities of detection" of at least 71.43% and "probabilities of false alarm" of at most 25.00%. The decision as to which feature and which threshold provide the best results again depends on the specifications which are desired. If a high "probability of detection" is desired, then the reflection coefficients with a threshold from 40% to 53% could be used, or the AR power spectral density with a threshold from 37% to 40% could be used. For these case, the "probability of detection" is 85.71% and the "probability of false alarm" is 25.00%. If a lower false alarm rate is desired, the AR power spectral density could be used with a threshold of 54% to 55%. For this case, the "probability of detection" is 71.43% and the "probability of false alarm" is 0.00%.

| Feature Vector | Minimum Percent SLS Threshold | Probability of Detection | Probability of False Alarm |
|-----------------------------|-------------------------------|--------------------------|----------------------------|
| AR power spectral density | 37.00 - 40.00 | 85.71 (6/7) | 25.00 (3/12) |
| AR power spectral density | 54.00 - 55.00 | 71.43 (5/7) | 0.00 (0/12) |
| ARMA power spectral density | 42.00 - 49.00 | 71.43 (5/7) | 16.67 (2/12) |
| reflection coefficients | 40.00 - 53.00 | 85.71 (6/7) | 25.00 (3/12) |

Table 3.4. "Probabilities of detection" and "probabilities of false alarm" which can be achieved from the middle band opening responses

3.3. High Bandwidth (7 kHz to 24 kHz) Opening Responses

The results obtained with the high bandwidth opening responses are given in Table 3.5. This table indicates the percentage of feature vectors for each sheep which were classified as SLS. These percentages are based upon the one-hundred opening responses which were extracted per sheep. Instances where at least one-half of the feature vectors were classified as SLS are indicated in bold.

For the sheep with intact valves, the percentage of feature vectors classified as SLS are again generally low. When the AR power spectral density was used as the feature vector, two sheep (75 and t520) had at least one-half of its vectors incorrectly classified as SLS. When the ARMA power spectral density was used as the feature vector only one sheep (75) had at least one-half of its vectors incorrectly classified as SLS. Lastly, when the reflection coefficients are used as the feature vector, only two sheep (414 and t520) had at least one-half of their vectors incorrectly classified as SLS.

For the sheep with SLS valves, the percentage of feature vectors classified as SLS are again generally higher than for the sheep with intact valves. When the AR power spectral density were used as the feature vector, five sheep (68, 103, 105, 412, and 422) had at least one-half of their vectors classified correctly as SLS. When the ARMA power spectral density was used as the feature vector, the percent of vectors classified correctly was not as high as those in Tables 3.1 and 3.3. For this case, only three sheep (68, 105, and 422) had at least one-half of their vectors classified correctly as SLS. Lastly, when the reflection coefficients were used, all seven sheep had at least one-half of their vectors classified correctly as SLS.

Fig's. 3.7 to 3.9 illustrate the "probabilities of detection" and the "probabilities of false alarm" which are achievable by varying the "minimum percent SLS threshold". Fig. 3.7 illustrates achievable probabilities using the AR power spectral density as the feature

| Sheep | Condition of Valve | Feature Vector | | |
|-------|--------------------------|---|---|--|
| | | AR power spectral density (120 features) (%) | ARMA power spectral density (120 features) (%) | Reflection coefficients (50 features) (%) |
| 40 | intact | 22.00 | 7.00 | 5.00 |
| 74 | intact | 26.00 | 35.00 | 0.00 |
| 75 | intact | 50.00 | 59.00 | 36.00 |
| 95 | intact | 1.00 | 2.00 | 15.00 |
| 129 | intact | 14.00 | 8.00 | 37.00 |
| 209 | intact | 10.00 | 11.00 | 16.00 |
| 414 | intact | 36.00 | 32.00 | 75.00 |
| 4337 | intact | 42.00 | 35.00 | 34.00 |
| 4380 | intact | 11.00 | 9.00 | 33.00 |
| bf1 | intact | 29.00 | 46.00 | 31.00 |
| bf12 | intact | 29.00 | 14.00 | 18.00 |
| t520 | intact | 80.00 | 18.00 | 54.00 |
| 68 | SLS | 80.00 | 51.00 | 83.00 |
| 103 | SLS | 58.00 | 44.00 | 55.00 |
| 105 | SLS | 60.00 | 53.00 | 65.00 |
| 110 | SLS | 31.00 | 41.00 | 68.00 |
| 411 | SLS | 34.00 | 27.00 | 57.00 |
| 412 | SLS | 63.00 | 23.00 | 91.00 |
| 422 | SLS | 50.00 | 50.00 | 80.00 |

Table 3.5. Percentage of features vectors which were classified as SLS for high band opening responses

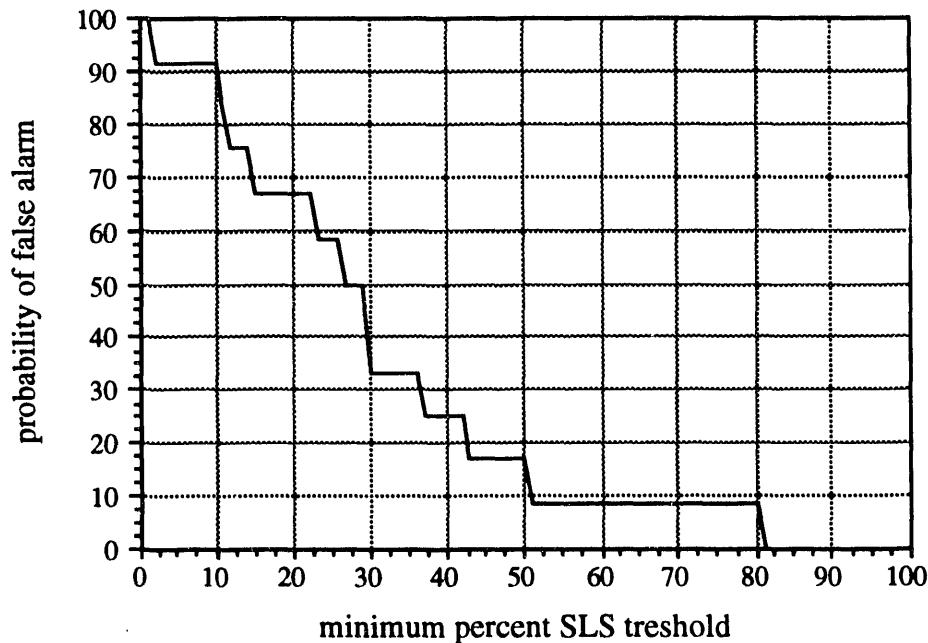
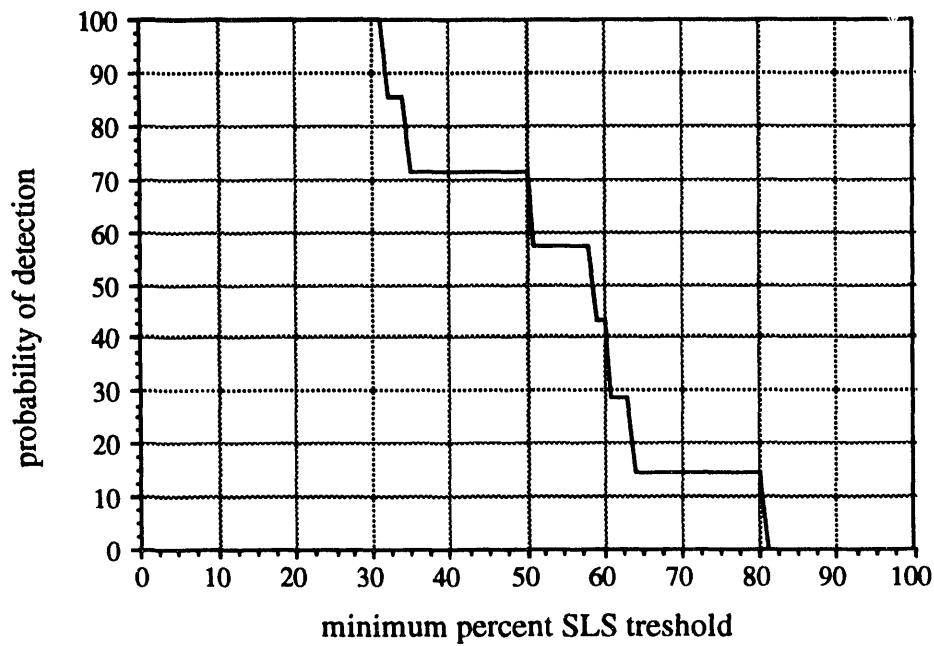


Figure 3.7. "Probability of detection" and "probability of false alarm" versus "minimum percent SLS threshold" for classification using AR power spectral density for high band openings

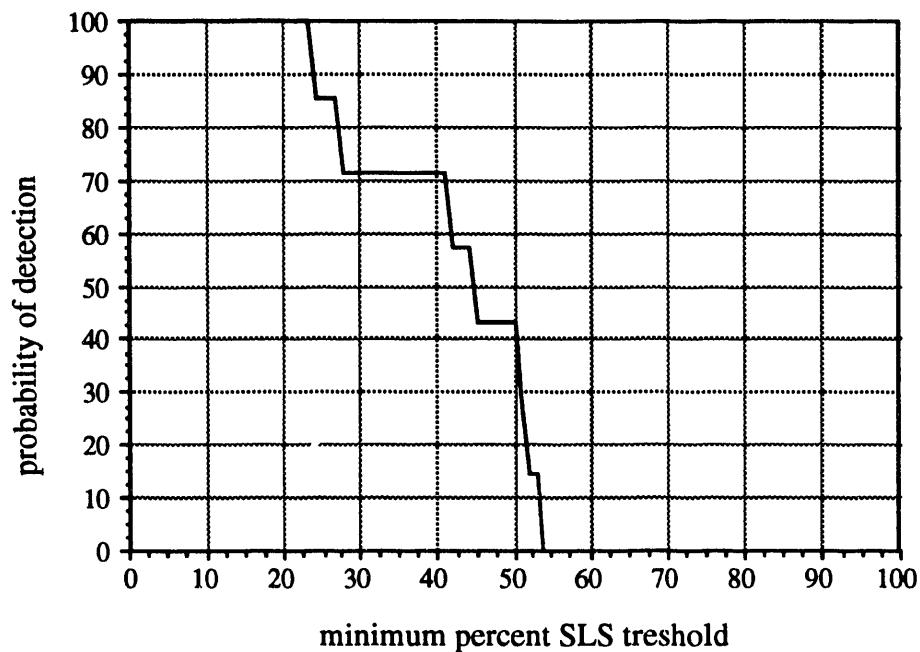


Figure 3.8. "Probability of detection" and "probability of false alarm" versus "minimum percent SLS threshold" for classification using ARMA power spectral density for high band openings

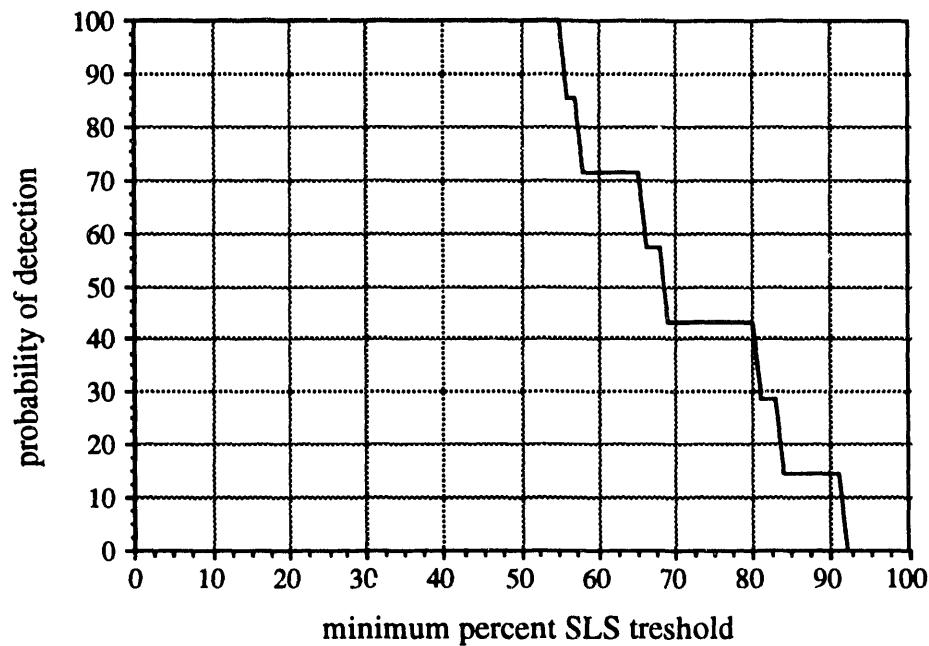


Figure 3.9. "Probability of detection" and "probability of false alarm" versus "minimum percent SLS threshold" for classification using reflection coefficients for high band openings

| Feature Vector | Minimum Percent SLS Threshold | Probability of Detection | Probability of False Alarm |
|------------------------------|-------------------------------|--------------------------|----------------------------|
| AR power spectral estimate | 43.00 - 50.00 | 71.43 (5/7) | 16.67 (2/12) |
| ARMA power spectral estimate | 36.00 - 41.00 | 71.43 (5/7) | 16.67 (2/12) |
| reflection coefficients | 55.00 | 100.00 (7/7) | 8.33 (1/12) |

Table 3.6. "Probabilities of detection" and "probabilities of false alarm" which can be achieved from the high band opening responses

vector, Fig. 3.8 illustrates achievable probabilities using the ARMA power spectral density as the feature vector, and Fig. 3.9 illustrates achievable probabilities using the reflection coefficients as the feature vector.

Table 3.6 present the best operating points from the graphs which result in "probabilities of detection" of at least 71.43% and "probabilities of false alarm" of at most 25%. From this table, it is apparent that the reflection coefficients with a threshold of 55% provided the best results. For this case, the "probability of detection" was 100.00% and the "probability of false alarm" was 8.33%.

4. Conclusions

Several conclusions can be drawn based on the results of this work. These are given below.

- "Single-leg separated" heart valves can be detected with "probabilities of detection" of at least 71.43% and corresponding "probabilities of false alarm" of at most 25.00% using any of the three bandwidths (1 kHz to 5 kHz, 3 kHz to 7 kHz, and 7 kHz to 24 kHz) of opening responses.

The probabilities which were achieved for all three bandwidths are shown in Tables 3.2, 3.4, and 3.6. When these tables are compared with each other, it is apparent that all bandwidths result in similar "probabilities of detection" and "probabilities of false alarm". These results were surprising because the signal-to-noise ratio of the responses in the high band were much greater than in the low and middle bands. Therefore, it was expected that the high band would provide better results. However, this was not the case. This may be due to the fact that the features which were extracted were all based on parametric models of the responses, and therefore the effect of noise may have been minimized. In any event, these results indicate that all bandwidths contain information which is useful in the detection process.

- "Single-leg separated" heart valves can be detected with "probabilities of detection" of at least 71.43% and corresponding "probabilities of false alarm" of at most 25.00% using features which represent the AR power spectral estimate, the ARMA power spectral estimate, and the reflection coefficients.

The probabilities which were achieved for all three feature vectors are shown in Tables 3.2, 3.4, and 3.6. When the results using these feature vectors are compared, it is apparent that they all yield similar results. The fact that the AR power spectral density and the reflection coefficients provide similar results was not surprising since both feature vectors are based upon all-pole models of the responses. The conversion from the n^{th} order AR coefficients and the n^{th} order reflection coefficients can be easily accomplished. On the other hand, the ARMA spectral density is based upon a pole-zero model of the response. Therefore, it was expected that it may produce slightly different results than the other two feature vectors. However, since the AR power spectral densities and the ARMA power spectral densities appeared similar and produce similar results, it appears as if the effects of the zeros in the model are insignificant. In any event, these results indicate that all the feature vectors contain information which is useful in the detection process.

- If one bandwidth and one feature vector had to be chosen based on the results of this work, the high bandwidth opening response with the reflection coefficients should be used with the nearest neighbor classifier.

For this case, a "probability of detection" of 100.00% and a "probability of false alarm" of 8.33% was achieved with a "minimum percent threshold" of 55.00%.

REFERENCES

- [1] "Artificial Heart Valves That Fail Are Linked To Falsified Records", *The Wall Street Journal*, Thursday, November 7, 1991.
- [2] Graham Thomas, "Signal Processing and Classification of Acoustic Signatures from Björk-Shiley Convexo-Concave Heart Valves" Lawrence Livermore National Laboratory internal document, September, 1992.
- [3] Carmen Mullenhoff, "Signal Processing of Shiley Heart Valve Data for Fracture Detection", UCRL-ID-113760, April 1, 1993.
- [4] R.O. Duda and P.E. Hart, *Pattern Classification and Scene Analysis*: Wiley, 1973.
- [5] J.V. Candy, *Signal Processing - The Modern Approach*. McGraw-Hill, 1988.
- [6] D. F. Specht, "Probabilistic Neural Networks," *Neural Networks*, Vol.3, pp109-118, 1990.

Appendix A

Feature Vector Description

A. Feature Vector Description

This appendix describes the features which were analyzed for the study presented in this report. These features are given below.

- autoregressive (AR) coefficients
- autoregressive (AR) power spectral densities
- auto-regressive, moving-average (ARMA) coefficients
- auto-regressive, moving-average (ARMA) power spectral densities
- reflection coefficients
- first-order statistics describing the shape of the ARMA power spectral densities
- first-order statistics describing the shape of the AR power spectral densities

Each one of these features represented one feature vector which was analyzed. These features are briefly described in sections A.1 to A.6. For a more thorough discussion refer to [5].

A.1. Autoregressive (AR) Coefficients

The autoregressive model is used to predict the responses of a signal based solely on past responses. Moreover, these estimates are represented as linear summations of the weighted past responses, where the weights are referred to as the autoregressive coefficients. This is represented mathematically as

$$\hat{y}(t) = \sum_{i=1}^{N_a} a_i y(t-i) \quad (A.1)$$

In this equation, $\hat{y}(t)$ is the estimate of the signal at time "t", a_i is the i^{th} autoregressive coefficient, and $y(t-i)$ is the signal which is being modeled at time "t - i". The order of the

model is the number of past responses which are used. In Eqn. (A.1), the order is N_a . The ability of this model to represent the signal is determined by the prediction error. This error, which is the difference between the estimated and actual signal, is given by

$$e(t) = y(t) - \hat{y}(t) = \sum_{i=0}^{N_a} a_i y(t-i) \quad , \text{ where } a_0 = 1. \quad (A.2)$$

An algorithm exists which determines the autoregressive coefficients in such a way so as to minimize the variance of the prediction error, which is defined as

$$R_{ee} = E\{e^2(t)\}. \quad (A.3)$$

This algorithm, which is referred to as the Levinson-Durbin algorithm, also computes this minimized error variance. Because the Levinson-Durbin algorithm computes these values based upon the optimization (minimization) of Eqn. (A.3), it is known as an optimal estimator. Optimal estimators which perform correctly have the important property of producing prediction error sequences which are zero-mean and white. These properties of the error will become important when the power spectral density is derived from the autoregressive model, as discussed in the next section.

A.2. Autoregressive (AR) Power Spectral Density

The autoregressive power spectral density is derived from the autoregressive model of a signal. However, before this is done, it will be helpful to first define the power spectral density. The power spectral density of a signal $y(t)$ represented in the z-domain is given by

$$S_{yy}(z) = E\{Y(z)Y^*(z)\}, \quad (A.4)$$

where $Y^*(z)$ is the complex conjugate of $Y(z)$. The first step in the derivation of this quantity is to define $Y(z)$ in terms of the autoregressive model. This is accomplished by first taking the z -transform of Eqn. (A.2). The expanded result of this is

$$E(z) = a_0 Y(z) + a_1 z^{-1} Y(z) + a_2 z^{-2} Y(z) + \dots + a_{N_a} z^{-N_a} Y(z). \quad (A.5)$$

Solving for $Y(z)$ and substituting $a_0 = 1$ yields

$$Y(z) = \frac{E(z)}{1 + a_1 z^{-1} + a_2 z^{-2} + \dots + a_{N_a} z^{-N_a}} = \frac{E(z)}{A(z)}. \quad (A.6)$$

Substituting Eqn. (A.6) into Eqn. (A.4) results in

$$S_{yy}(z) = E\{Y(z)Y^*(z)\} = \frac{1}{E\{A(z)A^*(z)\}} E\{E(z)E^*(z)\}. \quad (A.7)$$

This can be simplified to

$$S_{yy}(z) = \left| \frac{1}{A(z)} \right|^2 S_{ee}(z). \quad (A.8)$$

Since the error sequence is zero-mean and white,

$$S_{ee}(z) = R_{ee}, \quad (A.9)$$

Upon substitution of Eqn. (A.9) into Eqn. (A.8), the power spectral density of $y(t)$ is given as

$$S_{yy}(z) = \left| \frac{1}{A(z)} \right|^2 R_{ee}. \quad (A.10)$$

This result can be expressed in the frequency domain by evaluating z on the unit circle. This is accomplished by substituting $z = e^{j\Omega}$ into Eqn. (A.10). Since the error variance and the autoregressive coefficients are known from the Levinson-Durbin algorithm, the power spectral density of $y(t)$ can be computed using Eqn. (A.10).

One last point should be made regarding the autoregressive model. From Eqn. (A.6), it is evident that the signal $y(t)$ can be modeled as the output of a system driven by white noise. This is represented in the z-domain below.

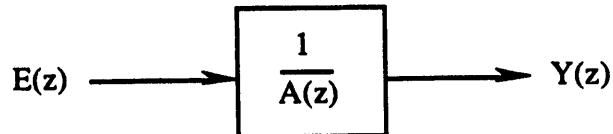


Figure A.1. Autoregressive model of a signal using z-transforms

Because this system only consists of a polynomial in the denominator, it is solely defined by the location of its poles. Therefore, the autoregressive model of a signal is often referred to as an all-pole model.

A.3. Autoregressive, Moving-Average (ARMA) Coefficients

The autoregressive, moving-average model of a signal is used to predict the response of a signal based on past responses and past prediction errors. These estimates are represented as a linear summation of the past responses and the past prediction errors. The weights on the past responses are referred to as the autoregressive coefficients, and the weights on the past prediction errors are referred to as the moving-average coefficients. This is represented mathematically as

$$\hat{y}(t) = \sum_{i=1}^{N_a} a_i y(t - i) + \sum_{j=1}^{N_c} c_j e(t - j) \quad (A.11)$$

In this equation $\hat{y}(t)$ is the estimate of the signal at time "t", a_i is the i^{th} autoregressive coefficient, $y(t - i)$ is the signal which is being modeled at time "t - i", c_j is the j^{th} moving-average coefficient, and $e(t - j)$ is the estimation error at time "t - j". The autoregressive order is the number of past responses which are used, and the moving-average order is the

number of past prediction errors which are used. In Eqn. (A.11), the autoregressive order is N_a , and the moving-average order is N_c . The error given in Eqn. (A.11) is defined as

$$e(t) = y(t) - \hat{y}(t) . \quad (A.12)$$

Substituting Eqn. (A.11) into Eqn. (A.12) results in an equation consisting only of the signal which is being modeled, the prediction errors, and the coefficients. This is given below

$$e(t) = y(t) - \hat{y}(t) = \sum_{i=0}^{N_a} a_i y(t-i) + \sum_{j=1}^{N_c} c_j e(t-j) , \text{ where } a_0 = 1. \quad (A.13)$$

The ability of this model to represent the signal is determined by the prediction error. One technique which determines the autoregressive and the moving-average coefficients in such a way so as to minimize this error variance is the recursive least-squares method [5]. This algorithm also computes the resulting error variance. Because these quantities are computed based upon the optimization (minimization) of Eqn. (A.3), this technique is an optimal estimator. Therefore, as discussed in section A.1, the prediction error sequence is zero-mean and white. These properties of the error will become important when the power spectral density is derived from the autoregressive, moving-average model, as discussed in the next section.

A.4. Autoregressive, Moving-Average (ARMA) Power Spectral Density

The autoregressive, moving-average power spectral density is derived from the autoregressive, moving-average model of a signal. The power spectral density of a signal $y(t)$ is given by Eqn. (A.4). As done in section A.2, the first step is to define $Y(z)$ in terms of the autoregressive model, moving-average model. This is accomplished by first taking the z-transform of Eqn. (A.13). This results in

$$E(z) = a_0 Y(z) + a_1 z^{-1} Y(z) + a_2 z^{-2} Y(z) + \dots + a_{N_a} z^{-N_a} Y(z) - c_1 E(z) + c_2 z^{-1} E(z) + c_3 z^{-2} E(z) + \dots + c_{N_c} z^{-N_c} E(z). \quad (A.14)$$

Solving for $Y(z)$ and substituting $a_0 = 1$ yields

$$Y(z) = \frac{1 + c_1 z^{-1} + c_2 z^{-2} + \dots + c_{N_c} z^{-N_c}}{1 + a_1 z^{-1} + a_2 z^{-2} + \dots + a_{N_a} z^{-N_a}} = \frac{C(z)}{A(z)} E(z). \quad (A.15)$$

Substituting Eqn. (A.15) into Eqn. (A.4) results in

$$S_{yy}(z) = E\{Y(z)Y^*(z)\} = \frac{E\{C(z)C^*(z)\}}{E\{A(z)A^*(z)\}} E\{E(z)E^*(z)\}. \quad (A.16)$$

This can be simplified to

$$S_{yy}(z) = \left| \frac{C(z)}{A(z)} \right|^2 S_{ee}(z). \quad (A.17)$$

Since the error sequence is zero-mean and white, its power spectral density is given by Eqn. (A.9). Upon substitution of Eqn. (A.9) into Eqn. (A.17), the power spectral density of $y(t)$ is given as

$$S_{yy}(z) = \left| \frac{C(z)}{A(z)} \right|^2 R_{ee}. \quad (A.18)$$

This result can be expressed in the frequency domain by evaluating z on the unit circle. This is accomplished by substituting $z = e^{j\Omega}$ into Eqn. (A.18). Since the error variance, the autoregressive coefficients, and the moving-average coefficients are known from the recursive least-squares algorithm, the power spectral density of $y(t)$ can be computed using Eqn. (A.18).

One last point should be made regarding the autoregressive, moving-average model. From Eqn. (A.15), it is evident that the signal $y(t)$ can be modeled as the output of a system driven by white noise. This is represented in the z -domain below.

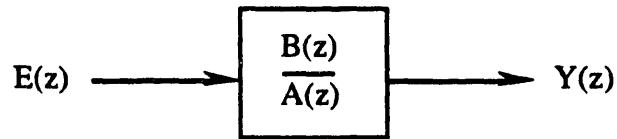


Figure A.2. Autoregressive, moving-average model of a signal using z-transforms

Because this system consists of polynomials in the denominator and the numerator, it is defined by poles and zeros. Therefore, the autoregressive, moving-average model of a signal is often referred to as a pole-zero model.

A.5. Reflection Coefficients

The reflection coefficients are used in conjunction with a lattice filter as a means of modeling a signal. The feed-forward, all-zero lattice structure is shown below.

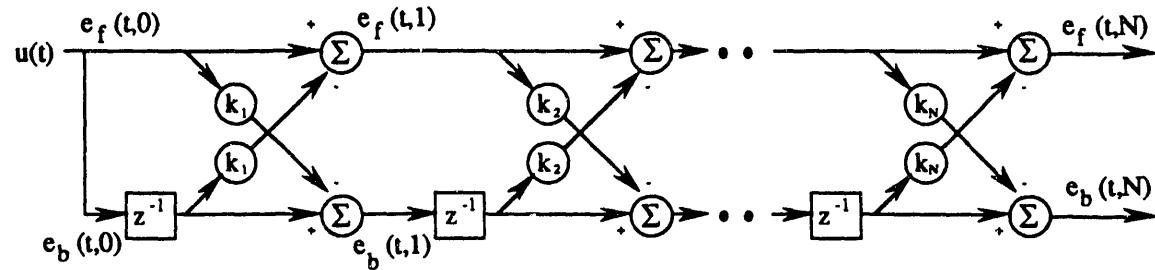


Figure A.3. Feed-forward lattice (all-zero) filter

The lattice model derives its physical origins from the fact that it mathematically represents a model of waves propagating through a layered medium. For instance, lattice models are employed to characterize seismic waves propagating through a tube in speech synthesis, or an electromagnetic wave propagating in a transmission line [5].

The lattice model is comprised of a number of stages. There are two inputs into each stage, which are referred to as the forward and backward prediction errors. The forward prediction error into stage "i" is represented by $e_f(t,i-1)$, and the backward prediction error into stage "i" is represented by $e_b(t,i-1)$. The outputs of stage "i" are also forward and

backward prediction errors which are represented by $e_f(t, i)$ and $e_b(t, i)$. The relationship between the input and output prediction errors for this stage are given by

$$\begin{aligned} e_f(t, i) &= e_f(t, i-1) - k_i e_b(t-1, i-1) \\ e_b(t, i) &= e_b(t-1, i-1) - k_i e_b(t, i-1), \end{aligned} \quad (A.19)$$

where k_i is the reflection coefficient for the stage. The order of the lattice model is represented by the number of stages. For the model in Fig. A.3, the order is N .

An algorithm exists which determines the reflection coefficients. This algorithm which is known as Burg's algorithm, assumes the signal which is to be modeled is input into the structure shown in Fig. A.3. The reflection coefficients are then computed so as to minimize the sum of the forward and backward error variance at each stage. This variance is given by

$$\text{Error variance} = E\{e_f^2(t) + e_b^2(t)\}. \quad (A.20)$$

Therefore, once the coefficients are determined, they can be used to characterize the signal.

A.6. First-Order Statistics Which Describe the Shape of the Power Spectral Densities

First-order statistics are often used to describe the shape of probability distributions. Therefore, it was thought that these statistics could be useful in characterizing the shape of a power spectrum. However, before these statistics could be computed, the power spectral density had to be normalized in such a way so that the sum of the densities for all spectral bins was one. This is shown by

$$\sum_{i=1}^N S_{yy}(f_i) = 1 \quad (A.21)$$

where N is the number of frequency bins. This result reflects the relative distribution of energy in the power spectral density. Once this normalization was accomplished, the first-order features could be computed. These features are given below.

$$\text{mean} = \mu = \sum_{i=1}^N f_i S_{yy}(f_i) \quad (A.22)$$

$$\text{standard deviation} = \sigma = \left[\sum_{i=1}^N (f_i - \mu)^2 S_{yy}(f_i) \right]^{0.5} \quad (A.23)$$

$$\text{skewness} = \frac{1}{\sigma^3} \sum_{i=1}^N (f_i - \mu)^3 S_{yy}(f_i) \quad (A.24)$$

$$\text{kurtosis} = \left[\frac{1}{\sigma^4} \sum_{i=1}^N (f_i - \mu)^4 S_{yy}(f_i) \right] - 3 \quad (A.25)$$

$$\text{energy} = \sum_{i=1}^N [S_{yy}(f_i)]^2 \quad (A.26)$$

$$\text{entropy} = - \sum_{i=1}^N S_{yy}(f_i) \log_2 [S_{yy}(f_i)] \quad (A.27)$$

**DATE
FILMED**

10/04/93

END

