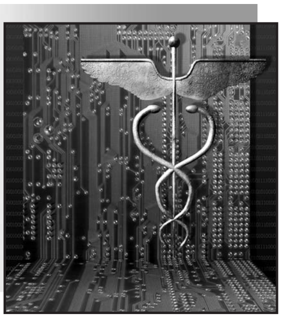

Automated Heart Abnormality Detection Using Sparse Linear Classifiers

A Computer-Aided Diagnosis System for Detecting Heart Wall-Motion Abnormalities from Echocardiograms

BY MALEEHA QAZI, GLENN FUNG, SRIRAM KRISHNAN, JINBO BI, R. BHARAT RAO, AND ALAN S. KATZ

Cardiovascular disease (CVD) is a global epidemic that is the leading cause of death worldwide (17 million deaths per year) [8]. It is the single largest contributor to "disability adjusted life years" (DALYs): 10% of DALYs in low- and middle-income nations and 18% of DALYs in high-income nations. Hence, the World Health Organization and the Centers for Disease Control agree that CVD is no longer an epidemic but a pandemic. In the United States, CVD accounted for 38% of all deaths in 2002 [7] and was the primary or contributing cause in 60% of all deaths. Coronary heart disease (CHD) accounts for more than half of CVD deaths (roughly 7.2 million deaths worldwide every year, and one of every five deaths in the United States), and it is the single largest killer in the world.

It is well known that early detection (along with prevention) is an excellent way of controlling CHD. CHD can be detected by measuring and scoring the regional and global motion of the left ventricle (LV) of the heart. It typically results in wall-motion abnormalities [i.e., local segments of the LV wall move abnormally (move weakly, not at all, or out of sync with the rest of the heart)], and sometimes motion in multiple regions, or indeed the entire heart, is compromised. The LV can be imaged in a number of ways. The most common method is the echocardiogram, which is an ultrasound video of different two-dimensional cross sections of the LV.

Unfortunately, echocardiograms are notoriously difficult to interpret, even for the best of physicians. Inter-observer studies have shown that even world-class experts agree on their diagnosis only 80% of the time [12], and intra-observer studies have shown a similar variation when the expert reads the same case twice at widely different points in time. There is a tremendous need for an automated "second-reader" system that can provide objective diagnostic assistance, particularly to the less-experienced cardiologist.

In this article, we address the task of building a computer-aided diagnosis system that can automatically detect wall-motion abnormalities from echocardiograms. We provide some medical background on cardiac ultrasound and the standard methodology used by cardiologists to score wall-motion abnormalities. We also describe our real-life dataset, which consists of echocardiograms used by cardiologists at St. Francis Heart Hospital to diagnose wall-motion abnormalities. We then provide an overview of our proposed system, which was built on top of an algorithm that detects and tracks the inner and outer walls of the heart [3]–[6]. It consists of a classifier that classifies the local region of the heart wall (and the entire heart) as normal or abnormal based on the wall motion. We also describe our methodology for feature selection and classification, followed by our experimental results.

Medical Background Knowledge

What Is Coronary Artery Disease?

The human heart is divided into four chambers: the left and right atrium and the left and right ventricle. The LV is the chamber responsible for pumping oxygenated blood to the entire body. As a result, it is the largest and strongest of the four chambers. Figure 1 shows the layout of the heart chambers in relation to one another; the LV is in the lower right part of the figure.

The heart is fed by three major coronary arteries: the left anterior descending (LAD), right coronary artery (RCA), and the left circumflex coronary artery (LCX). All three of these vessels feed the muscle surrounding the LV. Coronary artery disease results from the development of plaque within the artery, which usually deposits along the walls. When the plaque restricts normal blood flow to an extreme extent the patient will experience chest pain, known as angina. When the blood flow to the heart muscle is reduced, the function of that piece of muscle fed by the blocked artery will begin to become impaired. This is known as ischemia. This functional impairment can be seen from ultrasound images of the heart, also called echocardiograms (echos).

One of the first effects of coronary artery disease is that the motion of the heart wall during contraction will become impaired. Accurate regional wall-motion analysis of the LV is an essential component of interpreting echos to detect this effect.

Divisions of the Heart

There are many imaging modalities that have been used to measure myocardial perfusion, left ventricular function, and coronary anatomy for clinical management and research; for this project we chose to use echocardiography. The Cardiac Imaging Committee of the Council on Clinical Cardiology of the American Heart Association has created a standardized

 0739-5175/07/$25.00©2007IEEE

recommendation for the orientation of the heart, angle selection, and names for cardiac planes and number of myocardial segments [1]. This is the standardization used in this project. Echo images are collected from four standard views: apical 4 chamber (A4C), apical 2 chamber (A2C), parasternal long axis (PLAX) or apical 3 chamber (A3C), and parasternal short axis (PSAX) (shown in Figure 2). The planes used to cut the heart to display these standard views are displayed in Figure 3 from reference [2]. The long-axis view extends from the LV apex through the aortic valve plane. The short-axis view is perpendicular to the long-axis view resulting in a circular view of the LV. The four-chamber view is perpendicular to both the long- and short-axis views and includes the left and right ventricle and left and right atrium. If one rotates the 4-chamber view plane counterclockwise about 60°, the two-chamber view is obtained, which shows the LV and the left atrium.

The LV is divided into 17 myocardial segments. The short-axis view that results in a circular view of the LV can be taken at three locations: near the apex (apical), at the middle (mid-cavity), or near the base (basal). The most desirable view is the mid-cavity cut. If one lays these three resultant rings against one another, all segments of the heart are visible in relationship to one another, as shown in Figure 4 (modified from reference [1]). The LAD feeds segments 1, 2, 7, 8, 13, 14 and 17; the RCA feeds segments 3, 4, 9, 10 and 15; and the LCX feeds segments 5, 6, 11, 12, and 16.

Understanding the Data

The data are based on standard adult transthoracic B-mode ultrasound images collected from the four standard views described previously. Currently, we only utilize two of the four possible views: A4C and A2C, which show 12 of the 16 total segments [we ignore the apex (segment 17) since it is near impossible to measure]. These 12 views are enough to achieve our goal of classifying hearts. Even though we have images at different levels of stress (resting, low-dose stress, peak-dose

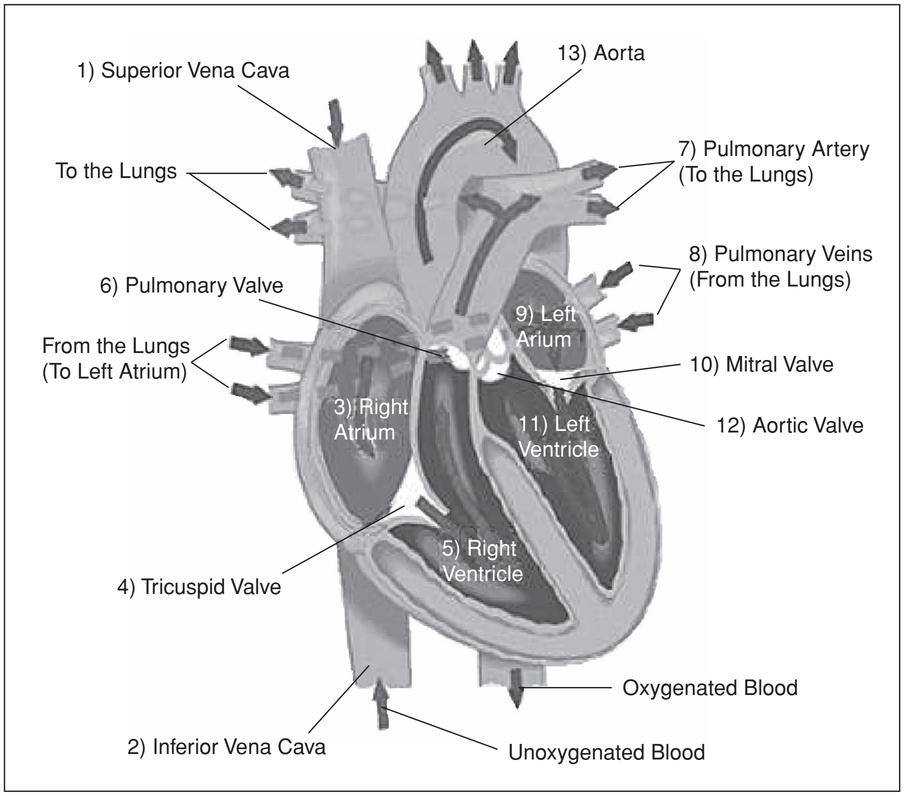

Fig. 1. Major parts of heart labeled, including the four chambers of the human heart: the left and right atrium, and the left and right ventricle.

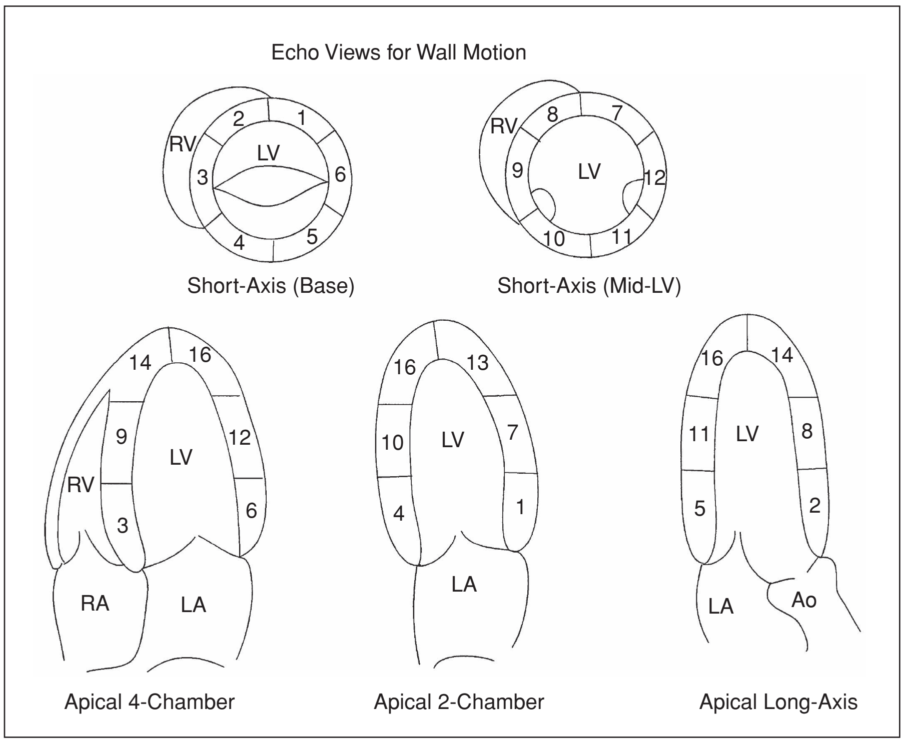

Fig. 2. Echocardiographic views for wall-motion evaluation. In the short-axis view, at the base and midventribular levels, the left ventricle is divided into anterior septum (2,8) and anterior free wall (1,7), lateral (6,12), posterior (5,11), inferior free wall (4,10), and posterior septal (3,9) segments. These same wall segments are seen in apical views as indicated, plus the anterior (13), septal (14), inferior (15), and lateral (16) apical segments are seen. Modified from reference (2) (segment numbers have been corrected to reflect standard naming convention being used).

stress, recovery), this work is based on images taken when the patient was resting. The goal is to automatically provide an initial score, or classification, to determine whether a heart is normal or abnormal given the ultrasound.

The echo data was collected from St. Francis Heart Hospital in Roslyn, New York. The data consist of 141 cases that will be used for training and 59 cases that are earmarked as the final test set; all of which were generated using exercise stress. All the cases have been labeled at the segment level by a group of trained cardiologists. The heart-level classification labels can be obtained from the segment-level labels by applying the following definition given to us by the doctors: a heart is considered abnormal if two or more segments are abnormal.

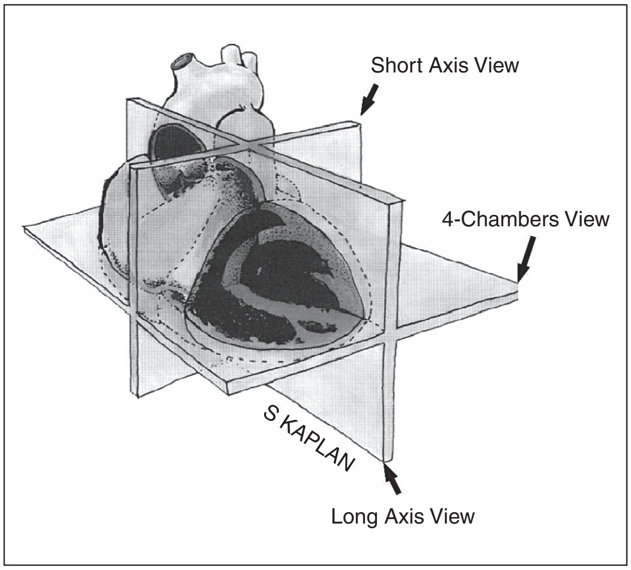

Fig. 3. The three basic image planes used in transthoracic echocardiography. The ventricles have been cut away to show how these image planes intersect the left and right ventricles. Dashed lines indicated the image planes at the great vessel and the atrial levels. From reference (2).

Preparation of the Data

Our application consists of two main parts: image processing and classification. The echos are run through an algorithm that automatically detects and tracks both the interior (endocardial) and exterior (epicardial) borders of the LV [4], [6]. Motion interferences (e.g., probe motion, patient movement, respiration, etc.) are compensated for by using global motion estimation based on robust statistics outside the LV. This is done so that only the heart's motion is analyzed. Then numerical feature vectors, which are extracted from the dual contours tracked through time, form the basis for the regional wall-motion classification.

Image Processing

The first step toward classification of the heart involves automatic contour generation of the LV [5]. Ultrasound is known to be noisier than other common medical imaging modalities such as MRI or CT, and echos are even worse due to the fast motion of the heart muscle and respiratory interferences. The framework used by our algorithm is ideal for tracking echo sequences since it exploits heteroscedastic (i.e., location-dependent and anisotropic) measurement uncertainties. The process can be divided into two steps: border detection and border tracking. Border detection involves localizing the LV on multiple frames of the image clip (shown in Figure 5 as a box drawn around the LV), and then detecting the LV's shape within that box. Seventeen control points are placed along the interior border of the LV to show where the border was detected. These points are then extended outward to find the external (epicardial) border of the LV.

Border tracking involves tracking both these contours together from one frame to the next through the entire movie clip. Motion interferences (e.g., probe motion, patient movement, respiration, etc.) are compensated for by using global motion estimation based on robust statistics outside the LV. This global motion estimation can be seen in Figure 6 as a vertical line near the center of the image.

After detection and tracking, numerical features are computed from the dual contours tracked through time. The features extracted are both global (involving the whole LV) and local (involving individual segments visible in the image) and are based on velocity, thickening, timing, volume changes, etc.

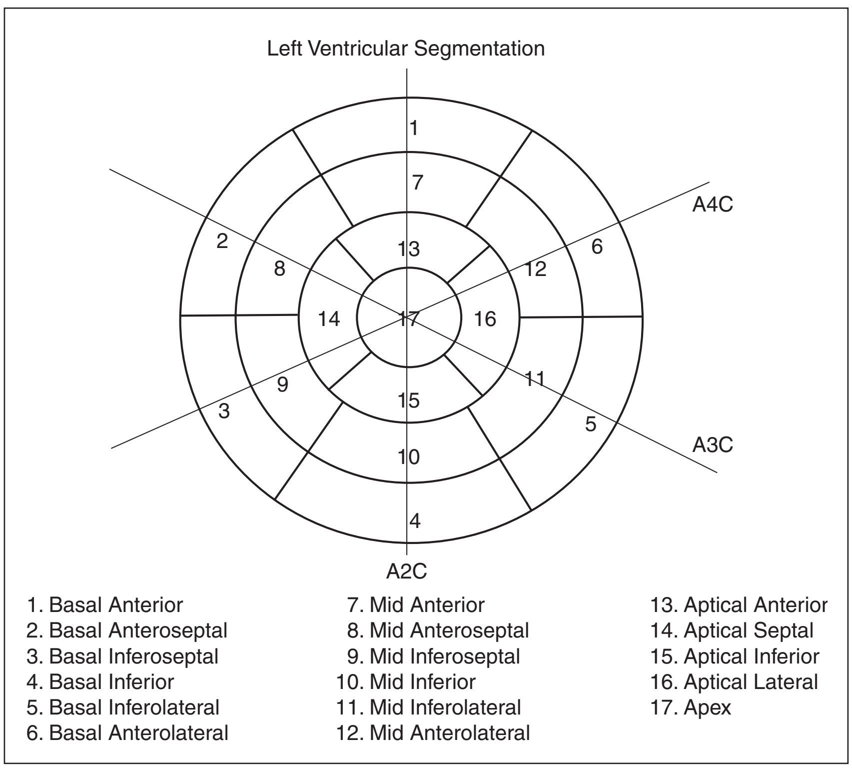

Fig. 4. Display, on a circumferential polar plot, of the 17 myocardial segments and the recommended nomenclature for tomographic imaging of the heart. Modified from reference (1).

Extracted Features

A number of features have been developed to characterize cardiac motion in order to detect wall-motion abnormalities, among them: global and local ejection fraction (EF) ratio, radial displacement, circumferential strain, velocity, thickness, thickening, timing, eigenmotion, curvature, and bending energy. Some of these features, including timing, eigenmotion, curvature, local EF ratio, and bending energy, are based only on the inner (endocardial) contour. Due to the patient examination protocol, only the systole (i.e., contraction phase of the heart) is recorded for some patients. In order for the features to be consistent, the systole is extracted from each patient based on the cavity area change. For each frame, the LV cavity area can be estimated accurately based on the inner (endocardial) contour of that frame. The frame corresponding to the maximal cavity area that is achieved at the end of diastolic phase (expansion phase of the heart) is the frame considered to be the beginning of systole. The frame corresponding to the minimal cavity area (achieved at the end of systolic phase) is the frame assumed to be the end of systole. For the time being, all features are computed based only on the systolic phase. However, the methods used to calculate the features are generally applicable for the diastolic phase as well.

The following list describes some of the many features.

- Timing-based features examine the synchronousness of the cardiac motion; i.e., whether all the points along the LV move consistently or not.
- Eigenmotion-based features determine the most significant moving direction of a point and the amount of its motion in that direction.
- Curvature-based features are mainly aimed at detecting abnormalities at the apex. This is also useful in identifying more general abnormalities associated with cardiac shapes. If a segment is dead, it may still move because it is connected to other segments, but we can observe that its shape will largely remain unchanged during the cardiac cycles. Curvature can capture this type of information.
- Local EF ratio features are aimed at capturing local cardiac contraction abnormalities.
- Bending energy features of the contour, assuming that the provided contour is made of elastic material and moving under tension, may be used to capture the cardiac contraction strength of a segment or the whole LV.
- Circumferential strain features, also called fractional shortening, measure how much the contour between any two control points shrinks in the systolic phase.

We had a total of 192 local (i.e., calculated per segment) and global (i.e., involving the whole LV, as shown in any one view) features, all of which were continuous. They included the features mentioned above as well as others not described here. As a general rule, the global version of certain features (e.g., radial displacement, radial velocity, etc) can be calculated by taking the mean, or standard deviation, of the six segments' respective feature values from any one view.

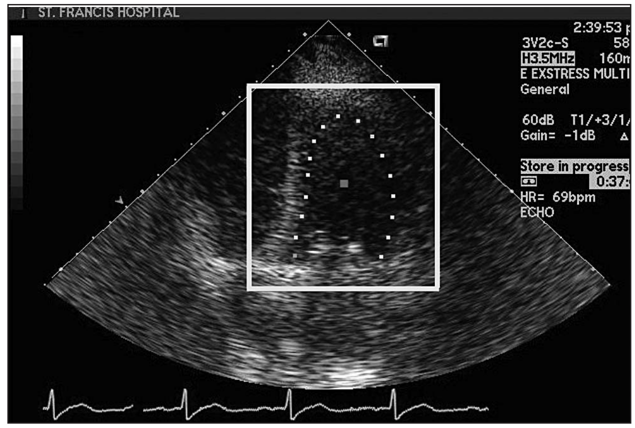

Fig. 5. One frame from an A4C image clip with the box showing the localized left ventricle, and the dots representing the control points along the detected inner contour.

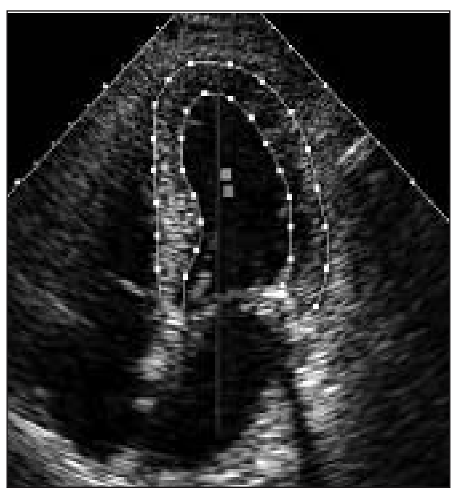

Fig. 6. One frame from an A4C image clip with the outer and inner contour control points shown. The vertical line near the middle shows use of global motion compensation, and the two squares denote the centers of the individual contours.

Data Mining

The classification algorithm used in the system is based on a novel feature selection technique, which is in turn based on mathematical programming. As a result, we obtain a hyperplane-based classifier that only depends on a subset of numerical features extracted from the dual contours tracked through time, and these are then used to provide classification for each segment and the entire heart.

Classification and Feature Selection

One of the difficulties in constructing a classifier for this task is the problem of feature selection. It is a well-known fact that reducing the feature dependence of a classifier improves the classifier's generalization capability. However, the problem of selecting an "optimal" minimum subset of features from a large pool (which is in the order of hundreds) of potential original features is known to be non-deterministic polynomial time (NP)-hard. Recently, Mika et al. proposed a novel mathematical programming formulation for linear Fisher's discriminant (LFD) using kernels [16], [15]. This new formulation included a regularization term similar to that used in the standard support vector machine (SVM) formulation [17]. We will make use of Mika's formulation but use a 1-norm instead of the 2-norm to obtain solutions that are more sparse and hence dependent on a smaller number of features. The next section describes the details of the approach.

Linear Fisher's Discriminant

The general idea behind LFD is to find the best subspace mapping such that it captures the best separation between

Echocardiograms are notoriously difficult to interpret, even for the best of physicians. Inter-observer studies have shown that even world-class experts agree on their diagnosis only 80% of the time.

the classes. Our problem involves binary classification; i.e., there are only two classes: positive (abnormal heart), and negative (normal heart) $\{\pm\}$.

Let $A_i \in R^{d\times l}$ be a matrix containing the l training data points on d-dimensional space and l_i the number of labeled samples for class i, $i \in \{\pm\}$. LFD [11] is the projection vector α, which maximizes,

$$J(\alpha) = \frac{\alpha^T S_B \alpha}{\alpha^T S_W \alpha} \tag{1}$$

where

$$S_B = (m_+ - m_-)\,(m_+ - m_-)^T$$

$$S_W = \sum_{i\in\{\pm\}} \frac{1}{l_i}\left(A_i - m_i e_{l_i}^T\right)\left(A_i - m_i e_{l_i}^T\right)$$

are the between and within class scatter matrices, respectively, and $m_i = (1/l_i)A_i e_{l_i}$ is the mean of class w_i and e_{l_i} is an l_i dimensional vector of ones. For (1) to be maximized, the numerator should be large, which represents the inter-class division (we want to push the classes as far apart as possible), and the denominator should be small, which represents the intra-class division (we want the points of any one class to be as near to one another as possible).

Transforming the above problem into a convex quadratic programming problem provides us some algorithmic advantages. First, notice that if α is a solution to (1), then so is any scalar multiple of it. Therefore, to avoid multiplicity of solutions, we impose the constraint $\alpha^T S_B \alpha = b^2$, which is equivalent to $\alpha^T(m_+ - m_-) = b$ where b is some arbitrary positive scalar. Then, the optimization problem (1) becomes

$$\begin{aligned} &\min_{\alpha\in R^d} && \alpha^T S_W \alpha \\ &s.t. && \\ & && \alpha^T(m_+ - m_-) = b \end{aligned} \tag{2}$$

For binary classification problems the solution of this problem is

$$\alpha^* = \frac{b S_W^{-1}\,(m_+ - m_-)}{(m_+ - m_-)^T S_W^{-1}\,(m_+ - m_-).} \tag{3}$$

According to this expansion, since S_W^{-1} is positive definite, unless the difference of the class means along a given feature is zero, all features contribute to the final discriminant. If a given feature in the training set is redundant, its contribution to the final discriminant would be artificial and not desirable. As a linear classifier, LFD is well suited to handle features of this sort provided that they do not dominate the feature set; i.e., the ratio of redundant to relevant features is not significant. Although the contribution of a single redundant feature to the final discriminant would be negligible when several of these features are available at the same time, the overall impact could be quite significant leading to poor prediction accuracy. Apart from this impact, in the context of LFD these undesirable features also pose numerical constraints on the computation of S_W^{-1} especially when the number of training samples is limited. Indeed, when the number of features, d is higher than the number of training samples, l, S_W becomes ill-conditioned and its inverse does not exist. Hence, eliminating the irrelevant and redundant features may provide a two-fold boost on the performance.

In what follows we propose a sparse formulation of LFD. The proposed approach incorporates a regularization constraint on the conventional algorithm and seeks to eliminate those features with limited impact on the objective function.

Sparse Linear Fisher's Discriminant Via Linear Programming

We propose a formulation similar to the one used for 1-norm SVM classifiers [9], where the 1-norm is introduced for both measuring the classification error and regulation. The use of the 1-norm instead of the 2-norm leads to linear programming formulations where very sparse solutions can be obtained. A sparse projection vector α implies that many input space features do not play a role in determining the linear classifier. In other words,

$$\alpha_i = 0 \Rightarrow \text{ the classifier does not depend on feature } i.$$

Our objective is to formulate an algorithm that can be seen as an approximation to (1) and that provides a sparse projection vector α. In order to achieve this, we add a regularization term to the objective function of (2):

$$\begin{aligned} &\min_{\alpha\in R^d} && \nu\alpha^T S_W \alpha + \|\alpha\|_1 \\ &s.t. && \\ & && \alpha^T(m_+ - m_-) = b \end{aligned} \tag{4}$$

where ν is the trade-off between $J(\alpha)$ maximization and regularization or sparsity of the projection vector α. The price to pay for sparsity of the solution is that, unlike (2), there is no a closed-form solution for the constrained quadratic in (4); furthermore, the parameter ν introduced in (4) has to be chosen by means of a tuning set that requires the problem to be solved several times and that can be computationally demanding. In order to address this issue we propose next a linear programming formulation that can be interpreted as an approximation to (4) and that results in sparser solutions than (4). Let's consider the following matrix:

$$H^T = \left[\frac{1}{\sqrt{l_+}} \left(A_+ - m_+ e_{l_+}^T\right)^T \quad \frac{1}{\sqrt{l_-}} \left(A_- - m_- e_{l_-}^T\right) \right].$$

From (1) we have that $S_w = H^T H$, then

$$\begin{aligned} \alpha^T S_W \alpha &= \alpha^T H^T H \alpha \\ &= (H\alpha)^T (H\alpha) \\ &= \|H\alpha\|_2^2. \end{aligned} \tag{5}$$

Hence the quadratic programming problem (4) can be rewritten as

$$\begin{aligned} \min_{\alpha \in R^d} \quad & \nu \|H\alpha\|_2^2 + \|\alpha\|_1 \\ s.t. \quad & \\ & \alpha^T (m_+ - m_-) = b. \end{aligned} \tag{6}$$

We can now use the 1-norm instead of the 2-norm in the objective function of (6) to obtain the following linear programming formulation that can be solved more efficiently and gives sparser solutions:

$$\begin{aligned} \min_{\alpha \in R^d} \quad & \nu \|H\alpha\|_1 + \|\alpha\|_1 \\ s.t. \quad & \\ & \alpha^T (m_+ - m_-) = b. \end{aligned} \tag{7}$$

That this problem is indeed a linear program can be easily seen from the equivalent formulation:

$$\begin{aligned} \min_{\alpha \in R^d} \quad & \nu e' s + e' t \\ s.t. \quad & \\ & \alpha^T (m_+ - m_-) = b \\ & -s \le H\alpha \le s \\ & -t \le \alpha \le t. \end{aligned} \tag{8}$$

Next, we propose an algorithm based on (8) and (3) that provides accurate LFD classifiers depending on a minimal set of features.

Algorithm 1 *Sparse Linear Fisher Discriminant*

Given the training dataset $\{A_-, A_+\}$ and a set of values $N = \{10^{-5}, 10^{-4}, \ldots, 10^5\}$ for the parameter ν do:

1) *For each $\nu \in N$ calculate cross-validation performance using the linear programming (8).*
2) *Let ν^* the value for which (8) gives the best cross-validation performance. Let's call $\hat{\alpha}$ the obtained sparse projection.*
3) *Select the subset $\hat{F}$ of the feature set F defined by $fi \in \hat{F} \Leftrightarrow \hat{\alpha}_i \neq 0$ or $f_i \in \hat{F} \Leftrightarrow |\hat{\alpha}_i| \geq tol$; that is, select the features corresponding to nonzero components of the projection $\hat{\alpha}$.*
4) *Solve original quadratic programming problem (1) with closed-form solution (3) considering only the feature subset $\hat{F}$ to obtain a final projection α^* that depends on only the "small" feature subset $\hat{F}$.*

Evaluation of Discovered Knowledge

Comparison to Other Algorithms

In order to empirically demonstrate the effectiveness of the proposed approach, we compared our feature-selection algorithm, sparse LFD (SLFD), to four other well-known classification algorithms. The first algorithm is a very popular publicly available implementation of SVM called SVMlight [13]. This formulation does not incorporate feature selection and produces classifiers that often depend on all the input features. The purpose of the comparison is to show that a feature-selection method improves generalization performance on this dataset. The second method included in our numerical comparisons is the relevance vector machine (RVM) algorithm [14], which is one of the most successful Bayesian methods for feature selection and sparse learning. It finds the relevance of features by optimizing the model marginal likelihood, also known as the evidence. The third approach consists of applying the standard LFD algorithm [11] without feature selection. The last classification approach used in our comparisons is the standard 1-norm SVM (SVM1) [9], which, similar to our approach, relies on the 1-norm regularization to obtain sparse classifiers. All the classifiers were trained using 141 cases and were tested on 59 cases. For the methods that needed parameters to be tuned (i.e., our algorithm and SVMlight), the model parameters were tuned by the means of leave-one-patient-out (LOPO) [10] cross validation on the training set. Ten-fold cross validation was not performed on this task because we wished to simulate a real-world situation where one does not have access to the test cases until the actual testing of the final classifier.

We have obtained many different answers from doctors as to what they feel the cost of a false positive (FP) (wrongly labeling the heart abnormal) or false negative (FN) (wrongly labeling the heart normal) happens to be. If this system is used as an initial reader, then too many FPs or FNs will cause the doctors to shut down the system because it is too unreliable. But as a validation system the main focus is to keep the FN rate low. In general, a high FP rate means you are sending too many patients for additional, more expensive tests, which would lead to higher costs for health insurance. A high FN rate could mean that a patient might go undiagnosed if the doctor using the system is not well trained and also misses potential abnormalities. For us, the "cost" of an FN is thus higher than an FP. By focusing on keeping the FN rate low, we lower the risk of missing abnormalities and leave the final diagnosis to the expertise of the doctor. Taking this into account, we decided that the best way to evaluate the classifier

Table 1. Areas under curve for the testing set and number of features selected for the five methods: SLFD, SVM_{light}, RVM, LFD, and SVM_1. (Best results shown in bold.)

Algorithm	*AUC*	# of features
SLFD	**89.6%**	**3**
SVM_{light}	87.4%	79*
RVM	85.8%	13
LFD	87.4%	79*
SVM_1	89.1%	8

*classifier uses all the features.

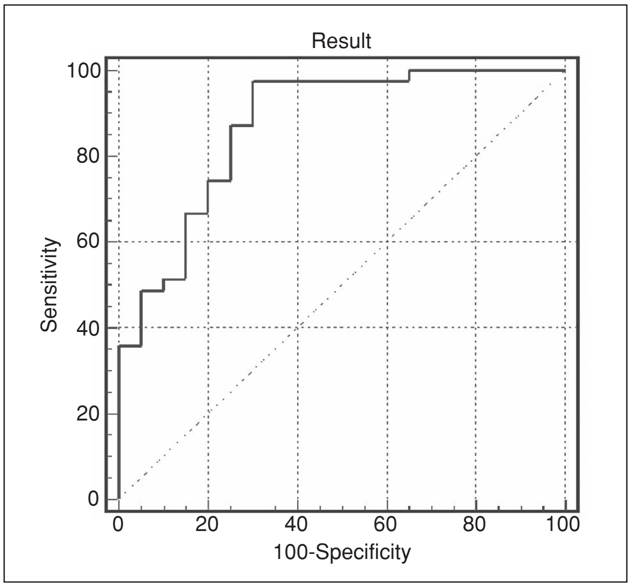

Fig. 7. ROC curve for the training set.

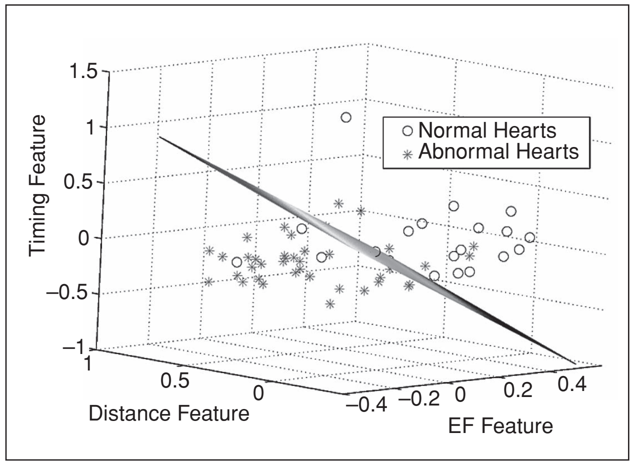

Fig. 8. Final hyperplane classifier in three dimensions: circles represent normal hearts and stars represent abnormal hearts in the test set.

performance is to measure the area under the curve (AUC) for receiver operating characteristic (ROC).

For each algorithm, Table 1 shows the AUC for the testing set and the number of features that the corresponding classifier depends on. As shown in the table, our method obtained the ROC with the largest area and depended on the fewest number of features (only three) of any of the algorithms tested. This low feature dependence is very important in our application since the features used for classification have to be calculated in real time.

Classification Results

The three features selected by SLFD were as follows:

- a feature that measures motion along the significant directions of movement of the walls of the heart
- a feature that measures correlation between the estimated area of the heart cavity and the distance between the walls of the heart to the center of mass of the heart
- the estimated EF of the heart.

It is important to note that two of the features (EF feature and the motion feature) were selected by all the classification methods tested. The performance obtained with SVM1 was the second best and was almost identical as the one obtained by SLFD but using eight features compared to only three used by SLFD. The ROC curve on the testing set for the final classifier is shown in Figure 7. The LOPO cross-validation performance for the final model was seven FPs and 17 FNs out of 81 positives (abnormals) and 60 negatives (normals); i.e., 88.3% of the normal hearts and 79.0% of the abnormal hearts were correctly classified.

On the testing set we obtained three FPs and 6 FNs out of 39 positives (abnormals) and 20 negatives (normals); i.e., 85.0% of the normal hearts and 84.6% of the abnormal hearts were correctly classified. A three-dimensional (3-D) plot depicting the final classifier and the test set is shown in Figure 8. These clinical results were presented and published at the American College of Cardiology meeting in March 2005 under the title "Clinical Evaluation of a Novel Automatic Real-Time Myocardial Tracking and Wall Motion Scoring Algorithm for Echocardiography Introduction."

Conclusion and Future Work

In this article we addressed the task of building an objective classification application for heart wall-motion analysis, based on features calculated off of echocardiograms. Our novel feature selection technique results in a robust hyperplane-based classifier that achieves the best performance in terms of AUC and number of features selected when compared to three other well-known classification algorithms.

The three features selected by our classifier (SLFD) are all global features, and their limited number makes it easier to explain the final classifier to physicians in order to get their feedback. In the future, we plan on expanding our classification to do segment-level classification for which we would identify different levels of CHD severity (normal, hypokinetic, akinetic, dyskinetic and aneurysm), incorporating the use of other standard echocardiography views (e.g., A3C, PSAX, PLAX) and including images from other levels of stress. We would also like to apply a ranking algorithm to take advantage of multiclass scores for classification. Comparisons of our proposed SLFD algorithm to other publicly available datasets and medical applications are also planned to further explore the potential of the algorithm.

Maleeha Qazi received a B.A. in mathematics from the University of Virginia in 1999 and an M.S. in computer science from the University of Wisconsin–Madison in 2002. She is currently working at Siemens Medical Solutions in Malvern, Pennsylvania. Her research interests include statistical relational learning, Bayesian networks, bioinformatics, and the application of machine learning techniques to the medical domain.

Glenn Fung received a B.S. in pure mathematics from Universidad Lisandro Alvarado in Barquisimeto, Venezuela, then earned an M.S. in applied mathematics from Universidad Simon Bolivar, Caracas, Venezuela, where later he worked as

an assistant professor for 2 years. Later, he earned an M.S. and a Ph.D. in computer sciences from the University of Wisconsin–Madison. His main interests are optimization approaches to machine learning and data mining, with emphasis on support vector machines. In the summer of 2003 he joined the Computer-Aided Diagnosis Group at Siemens Medical Solutions in Malvern, Pennsylvania, where he has been applying machine learning techniques to solve challenging problems that arise in the medical domain. His recent papers are available at www.cs.wisc.edu/~gfung.

Sriram Krishnan received the Ph.D. in electrical engineering (systems) from the University of Michigan in 1998. From 1998–2002, he worked as a staff researcher at Acuson Corporation, focusing on advanced ultrasound beamforming. In 2002, he joined the Computer-Aided Diagnosis Group at Siemens Medical Solutions. His research interests include application of decision support systems in cardiology and application of computer vision and machine learning for cardiac imaging.

Jinbo Bi received the B.S. in mathematics and M.S. in automatic control from Beijing Institute of Technology, Beijing, China, in 1995 and 1998, and the Ph.D. in mathematics from Rensselaer Polytechnic Institute, Troy, New York, in 2003. She is currently a staff scientist working with Siemens Medical Solutions. Her research interests include mathematical programming, machine learning, kernel methods, support vector machines, feature selection, drug discovery, and medical imaging.

R. Bharat Rao is the senior director of engineering R&D for the Computer-Aided Diagnosis and Therapy Solutions Group at Siemens Medical Solutions, Malvern, Pennsylvania. He received the B.Tech in electronics engineering from the Indian Institute of Technology, Madras, in 1985, and the Ph.D. in electrical engineering from the University of Illinois, Urbana-Champaign, in 1993. His research interests include machine learning, classification, graphical models, probabilistic inference, and optimization, with a focus on developing decision-support systems that can help physicians improve the quality of patient care. He is particularly interested in the development of novel data mining methods to collectively mine and integrate the various parts of a patient record (lab tests, pharmacy, free text, images, proteomics, etc.) and the integration of medical domain knowledge into the mining process. He received the "Siemens Inventor of the Year" award in 2005.

Alan S. Katz received the B.S. in biology from Brandeis University in 1978 and the M.S. in engineering science/bioinformation systems from the California Institute of Technology in 1980. In 1984, he received his M.D. from the

University of Vermont. Dr. Katz did his residency in internal medicine at North Shore/ Memorial-Sloan Kettering and his fellowship in cardiology at New York Hospital-Cornell Medical Center. He served as director of echocardiography at Miriam Hospital, Rhode Island, and was on the faculty of Brown University from 1987 to 2000. He is currently director of cardiac imaging and informatics at St. Francis Hospital—The Heart Center in Roslyn, New York. He holds appointments as associate professor of medicine at SUNY Stony Brook and adjunct associate professor of medicine at Brown. Dr. Katz's research interests include echocardiography and digital imaging in medicine.

Address for Correspondence: Glenn Fung, 4909 Woodburn Drive, Madison, WI 53711. E-mail: glenn.fung@siemens.com, maleeha.qazi@siemens.com.

References

[1] M.D. Cerqueira, N.J. Weissman, V. Dilsizian, A.K. Jacobs, S. Kaul, W.K. Laskey, D.J. Pennell, J.A. Rumberger, T. Ryan, and M.S. Verani, "Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart: A statement for healthcare professionals from the Cardiac Imaging Committee of the Council on Clinical Cardiology of the American Heart Association," *Circulation*, vol. 105, pp. 539–542, Jan. 2002.

[2] C.M. Otto, *Textbook of Clinical Echocardiography*. 2nd ed. Philadelphia, PA: W.B. Saunders, 2000.

[3] X.S. Zhou, D. Comaniciu, and A. Gupta, "An information fusion framework for robust shape tracking," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 27, pp. 115–129, Jan. 2005.

[4] D. Comaniciu, "Nonparametric information fusion for motion estimation," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, Madison, WI, June 2003, vol. 1, pp. 59–66.

[5] B. Georgescu, X.S. Zhou, D. Comaniciu, and A. Gupta, "Database-guided segmentation of anatomical structures with complex appearance," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, San Diego, CA, 2005, pp. 429-436.

[6] D. Comaniciu, X.S. Zhou, and S. Krishnan, "Robust real-time tracking of myocardial border: An information fusion approach," *IEEE Trans. Med. Imag.*, vol. 23, pp. 849–860, July 2004.

[7] American Heart Association (2005), Heart disease and stroke statistics 2005 update [Online]. Available: http://www.americanheart.org/downloadable/heart.

[8] World Health Organization (2004), The atlas of global heart disease and stroke [Online]. Available: http://www.who.int/cardiovasculardiseases/resources/atlas/.

[9] P.S. Bradley and O.L. Mangasarian, "Feature selection via concave minimization and support vector machines," in *Proc. 15th Int. Conf. Machine Learning*, Madison, WI, July 1998, pp. 82–90.

[10] M. Dundar, G. Fung, L. Bogoni, M. Macari, A. Megibow, and B. Rao, "A methodology for training and validating a CAD system and potential pitfalls," in *Proc. Computer Assisted Radiology and Surgery Conf.*, Chicago, IL, June 2004, pp. 837–844.

[11] R. Duda, P. Hart, and D. Stork, *Pattern Classification*. New York: Wiley, 2001.

[12] R. Hoffmann, T.H. Marwick, D. Poldermans, H. Lethen, R. Ciani, P. van der Meer, H.-P. Tries, P. Gianfagna, P. Fioretti, J.J. Bax, M.A. Katz, R. Erbel, and P. Hanrath, "Refinements in stress echocardiographic techniques improve inter-institutional agreement in interpretation of Dobutamine stress echocardiograms," *Eur. Heart J.*, vol. 23, no. 10, pp. 821–829, May 2002.

[13] T. Joachims (2002), SVMlight [Online]. Available: http://svmlight.joachims.org.

[14] Y. Qi, T.P. Minka, R.W. Picard, and Z. Ghahramani, "Predictive automatic relevance determination by expectation propagation," in *Proc. 21st Int. Conf. Machine Learning*, Alberta, Canada, 2004, pp. 85-92.

[15] S. Mika, G. Ratsch, and K.-R. Muller, "A mathematical programming approach to the kernel Fisher algorithm," in *Proc. Neural Information Processing Systems 2000*, Denver, CO, pp. 591–597, 2000.

[16] S. Mika, G. Ratsch, J. Weston, B. Scholkopf, and K.-R. Muller, "Fisher discriminant analysis with kernels," in *Neural Networks for Signal Processing IX, Proc. 1999 IEEE Signal Processing Workshop*, Madison, WI, pp. 41–48, Aug. 1999.

[17] O.L. *Mangasarian*, "Generalized support vector machines," in *Advances in Large Margin Classifiers*, A. Smola, P. Bartlett, B. Scholkopf, and D. Schuurmans, Eds. Cambridge, MA: MIT Press, 2000, pp. 135–146.