

# Cardiac Disease Recognition in Echocardiograms Using Spatio-temporal Statistical Models

David Beymer and Tanveer Syeda-Mahmood

**Abstract**—In this paper we present a method of automatic disease recognition by using statistical spatio-temporal disease models in cardiac echo videos. Starting from echo videos of known viewpoints as training data, we form a statistical model of shape and motion information within a cardiac cycle for each disease. Specifically, an active shape model (ASM) is used to model shape and texture information in an echo frame. The motion information derived by tracking ASMs through a heart cycle is then represented compactly using eigen-motion features to constitute a joint spatio-temporal statistical model per disease class and observation viewpoint. Each of these models is then fit to a new cardiac echo video of an unknown disease, and the best fitting model is used to label the disease class. Results are presented that show the method can discriminate patients with hypokinesia from normal patients.

## I. INTRODUCTION

Echocardiography is often used to diagnose cardiac diseases related to regional motion, septal wall motion, as well as valvular motion abnormalities. It provides images of cardiac structures and their movements, giving detailed anatomical and functional information about the heart from several standard viewpoints such as apical 4-chamber view, parasternal long-axis view, etc. Cardiologists routinely use the echo videos to aid their diagnosis of the underlying disease. However, since the available tools provide few quantitative measurements about the complex spatio-temporal motion of the heart, physicians usually resort to ‘eye balling’ the disease. Discerning motion abnormalities such as when the myocardium contracts significantly less than the rest of the tissue, is difficult, in general for humans. Unlike the interpretation of static images such as X-rays, it is difficult for physicians to describe the nature of the abnormalities in moving tissues. Thus, tools that automate the disease discrimination process by capturing and quantifying the complex 3D non-rigid spatio-temporal heart motion can significantly aid the decision support process of physicians.

In this paper we present a method of automatic disease recognition using statistical spatio-temporal models of diseases in cardiac echo videos. Unlike previous approaches to cardiac disease discrimination that perform feature extraction and similarity search on global average motion [10], we perform a detailed modeling of the appearance (shape and texture) using the well-known Active Shape Models[3]. The motion information derived by tracking ASMs through a heart cycle is then represented compactly using eigen-motion features per disease class. Each of these models is then

fit to a new cardiac echo video of an unknown disease, and the best fitting model is used to recognize the disease class. Since different diseases are best diagnosed in different characteristic views, we build the joint spatio-temporal model for chosen viewpoints per disease class. The average fit for all such viewpoints is then used to rank the overall disease model fits. We demonstrate the effectiveness of the model fitting approach to disease discrimination on patients with hypokinesia, and show that they can be discriminated from normal patients.

The rest of the paper describes this approach in detail. In Section II, we review related work on automatic cardiac disease discrimination and spatio-temporal modeling of cardiac echo videos. In Section III, we describe our overall approach to disease recognition from spatio-temporal models. Finally, in Section IV, we present results of disease discrimination of hypokinesia patients over normal patients.

## II. RELATED WORK

The estimation of cardiac motion and deformation from cardiac imaging has been an area of major concentration in medical image analysis [4]. More work has been done on segmenting and motion tracking of heart regions such as left ventricle than classification of normal and abnormal motion patterns. These include methods for ventricular segmentation [5], motion tracking [7], measurements on tagged MRI [12] and functional measurements of blood flow [6]. Regional segmentation has been an active research area with all methods requiring some assistance in identifying initial candidate regions. Methods relying on phase contrast MRI also exist to study cardiac motion and deformation [6], [8]. In general, all these methods depend on an accurate segmentation of the myocardium walls, but have the advantages of being imaging modality independent. The dependency on obtaining an accurate segmentation, however, remains a significant issue, as there still are no fully automated robust and efficient left ventricle (LV) surface segmentation methods.

Recently, there have been efforts to discriminate diseases by analyzing spatio-temporal properties of heart regions. In [11], the left ventricular region was modeled using volumetric meshes to measure both geometric and motion information throughout a heart cycle. In [10], a global estimation of average direction and extent of motion was done using average velocity curves. Features were extracted from average velocity curves and clustered to discriminate various cardiac diseases. The global motion curves do not accurately capture the regional motion patterns observed in cardiac motion where the septal wall moves quite differently from the

valvular regions. Later work combined region segmentation using a graph-theoretic approach based on normalized cut [9] with motion signatures to improve disease discrimination. Even so, the inaccuracies of region segmentation using a graph-theoretic approach and rough motion estimation using average velocities often lead to inaccuracies in disease discrimination.

Finally, the approach presented here utilizes the machinery of Active Shape Models (ASMs) [3] and low-dimensional statistical models for motion. ASMs are an effective statistical tool for modeling the appearance and texture of nonrigid objects, and they have been extensively applied in computer vision. They have also been previously used for echocardiogram segmentation [1], tracking and for region localization and shape correspondence [9]. When combined with motion modeling, these statistical models present an alternate approach to disease discrimination based on precise modeling of spatio-temporal information. Specifically, by learning such models from several example echo videos per cardiac disease, we can express disease recognition as a problem of fitting the appropriate spatio-temporal model.

### III. SPATIO-TEMPORAL STATISTICAL MODELS

Our approach to disease discrimination is based on building spatial and temporal models from sample learning data for both disease and normal patient groups and using a matching algorithm for resolving the best matching model. Since different diagnostic viewpoints result in very different spatio-temporal sequences, it will be necessary to build spatio-temporal models separately per viewpoint. Further, since the viewpoints used to diagnose vary from disease to disease, the number of viewpoint-dependent spatio-temporal models will vary across diseases. For disease discrimination, therefore, we use spatio-temporal models from corresponding viewpoints for the purpose of model fitting. The overall approach is as follows. For each disease and each of its diagnostic viewpoints, we collect a set of patient training cardiac sequences  $Seq_{view}$  (Fig. 1, left) and we train two models: a spatial active shape model  $ASM_{view}$  to represent the spatial information and a motion model  $Mot_{view}$  to represent the motion within a heart cycle. The ASM models build a disease-specific model of cardiac appearance for individual frames. The motion model represents an entire sequence of tracked features by projecting them into a set of disease-specific “eigenmotions.” Given testing sequences from a new patient  $P$  with unknown cardiac disease, we select candidate ASM and motion models for the corresponding viewpoints from the learned disease models and try to fit the models the new sequences. For a given disease, the model fitting scores for the viewpoints are averaged to ensure a good fit to all viewpoint sequences taken from patient  $P$ . The disease label of the spatio-temporal model that shows the best average fit is then taken as the disease label for the new patient. Thus we combine evidence from different echocardiographic viewpoints and diseases using a common statistical framework.

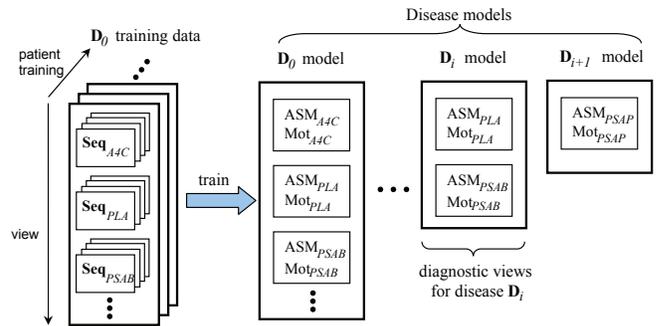


Fig. 1. Our system trains disease-specific models of echocardiogram appearance and motion for the diagnostically-important viewpoints for each disease. The models consist of an active shape model (ASM) for static appearance, and a motion model for tracked feature motion in the cardiac cycle. A4C, PLA, PSAB, and PSAP are acronyms for cardiac echo viewpoints, and stand for, respectively apical 4 chamber, parasternal long axis, parasternal short axis - basal level, and parasternal short axis - papillary muscle level.

#### A. Appearance modeling

Using the paradigm of Active Shape Models [3], we represent shape information in an echo frame by a set of feature points obtained from important anatomical regions depicted in the echo sequences. For example, in apical 4-chamber views, the feature points are centered around the mitral valve, interventricular septum, interatrial septum, and the lateral wall (Fig. 2). We chose not to trace the entire boundary of the left ventricle because the apex is often not visible due to the ultrasound zoom and transducer location. Wall boundaries are annotated on both the inner and outer sections. While the feature points are manually isolated during training stage, they are automatically identified during matching. Thus, the cardiac region information used during model building captures features important for cardiac assessment.

For each disease, a number of training images were collected from cardiac cycle video clips – covering different patients, diagnostically important views, and time offsets within the cycle. We assume that each echo video clip represents a full heart cycle. If more than one heart cycle is present, they can be segmented into cycles based on the synchronizing ECG often present in cardiac echo videos.

To model the appearance, we represent the features extracted from each echo video frame by a shape vector  $\mathbf{s}$  obtained by concatenating the locations  $(x, y)$  of  $n$  features  $f_1, f_2, \dots, f_n$  as

$$\mathbf{s} = [x_1, y_1, \dots, x_n, y_n]^T. \quad (1)$$

Next, we form image patches centered around the feature locations to capture texture information. Given a training image with  $n$  features, the texture vector  $\mathbf{t}$  concatenates the pixel intensities or RGB values from all the patches into one long vector, where patch size is matched to the pixel spacing between features. Fig. 2 shows the feature points we use for apical 4 chamber views. After shape annotation, the shape vectors  $\mathbf{s}$  are geometrically normalized by a similarity transform for model generation. The texture vector  $\mathbf{t}_{raw}$  is

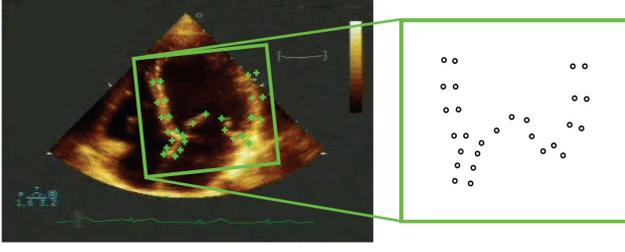


Fig. 2. ASM feature points for apical 4 chamber view. (Best viewed in color.)

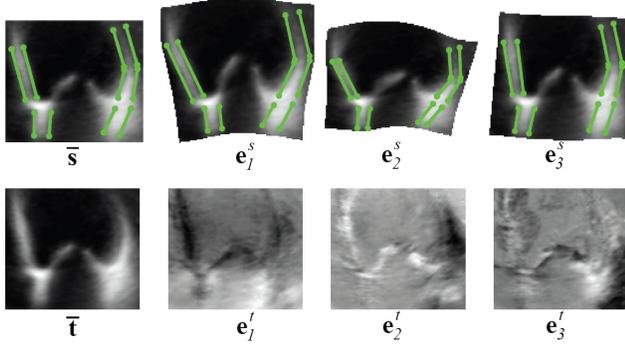


Fig. 3. Top row: Mean shape  $\bar{\mathbf{s}}$  and top shape eigenvectors. Bottom row: Mean texture  $\bar{\mathbf{t}}$  and top texture eigenvectors (for clarity of presentation, texture is shown as one image instead of patches).

similarly normalized for echo gain and offset by subtracting the mean and dividing by standard deviation as

$$\mathbf{t} = (\mathbf{t}_{raw} - \bar{\mathbf{t}}_{raw}) / \sigma(\mathbf{t}_{raw}). \quad (2)$$

To construct the ASM model, we reduce the dimensionality of the spatial and textural vector using PCA to find a small set of *eigenshapes* and *eigentextures*. The shape  $\mathbf{s}$  and texture  $\mathbf{t}$  can then be linearly modeled to form the active shape model as:

$$\begin{aligned} \mathbf{s} &= S\mathbf{a} + \bar{\mathbf{s}} & S &= \begin{bmatrix} | & | & \dots & | \\ \mathbf{e}_1^s & \mathbf{e}_2^s & \dots & \mathbf{e}_p^s \\ | & | & \dots & | \end{bmatrix} \\ \mathbf{t} &= T\mathbf{b} + \bar{\mathbf{t}} & T &= \begin{bmatrix} | & | & \dots & | \\ \mathbf{e}_1^t & \mathbf{e}_2^t & \dots & \mathbf{e}_q^t \\ | & | & \dots & | \end{bmatrix} \end{aligned} \quad (3)$$

where  $p$  eigenshapes  $\mathbf{e}_1^s, \mathbf{e}_2^s, \dots, \mathbf{e}_p^s$  and  $q$  eigentextures  $\mathbf{e}_1^t, \mathbf{e}_2^t, \dots, \mathbf{e}_q^t$  are retained in the PCA (see Fig. 3). The  $p$ -dimensional vector  $\mathbf{a}$  and the  $q$ -dimensional vector  $\mathbf{b}$  are the low dimensional representations of shape and texture.

### B. ASM Fitting

Fitting an ASM to a new image involves finding a similarity transform  $\Gamma^{\text{sim}}$  to position the model appropriately in the image and recovering the shape and texture vectors  $\mathbf{a}$  and  $\mathbf{b}$ . This is iteratively estimated by alternating between shape and texture update steps. Unlike most ASM techniques, which require a manual initialization to start the model fitting, we use an automatic initialization method where a

distance-to-eigenspace method is first used to generate seed ASM initializations. The seed locations are then refined and evaluated using a coarse-to-fine pyramid approach, where seed hypotheses are pruned if their fitting error rises above a threshold. The ASM seed hypotheses that reach the finest pyramid level in the first frame are then tracked across the rest of the frames of the sequence, with the fitting result for frame  $t$  serving as the initial conditions for frame  $t+1$ . If multiple hypotheses survive sequence tracking, they are ranked by their fitting score.

To evaluate an ASM fit at a given position, we measure the error of fit in shape space and texture space using the Mahalanobis distance and the texture reconstruction error suitably normalized. For image  $I$ , the ASM fit  $(\mathbf{a}, \mathbf{b}, \Gamma^{\text{sim}})$  is

$$\text{fit}(\mathbf{a}, \mathbf{b}, \Gamma^{\text{sim}}) = \mathbf{a}^T \Sigma_{\text{shp}}^{-1} \mathbf{a} + \mathbf{b}^T \Sigma_{\text{tex}}^{-1} \mathbf{b} + 2R^2 / \lambda_{\text{tex}}^{q+1}, \quad (4)$$

where  $R = \|\mathbf{t} - T\Gamma^{\text{sim}}\mathbf{t}\|$ ,  $\mathbf{t} = I(\Gamma^{\text{sim}}(x, y))$  is the extracted object texture from the input image,  $\lambda_{\text{tex}}^{q+1}$  is the  $(q+1)$ th texture eigenvalue,  $\Sigma_{\text{shp}}$  is a diagonal  $(p \times p)$  matrix with shape PCA eigenvalues, and  $\Sigma_{\text{tex}}$  is a diagonal  $(q \times q)$  matrix with texture PCA eigenvalues (see Cootes and Taylor [2]).

### C. Motion Modeling

Next, we consider the modeling of heart motion within a cardiac cycle. Using an ASM to track cardiac motion through a cycle with  $m$  frames will produce  $m$  sets of feature points  $\{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_m\}$ , where each  $\mathbf{s}_i$  is a column vector of stacked ASM  $(x, y)$  feature coordinates. To create a canonical representation, we vectorize the motion in the cardiac cycle by normalizing for image plane geometry and standardizing the time axis to a target length  $n$  as follows:

- 1) *Geometry normalization*. Align the first frame  $\mathbf{s}_1$  to a canonical position using a similarity transform  $\Gamma_1^{\text{sim}}$ . Apply the same similarity transform to standardize all frames

$$\mathbf{s}_i \leftarrow \Gamma_1^{\text{sim}}(\mathbf{s}_i).$$

Thus, all motion will be described in a coordinate frame that factors out scale, rotation, and translation.

- 2) *Time normalization*. Interpolate the  $\mathbf{s}_i$ 's in time using piecewise linear interpolation to standardize the sequence length from input length  $m$  to a target length  $n$ .
- 3) *Shape decoupling*. Factor out object shape by subtracting out frame 1, creating our final motion vector  $\mathbf{m}$

$$\mathbf{m} = [\mathbf{s}_2 - \mathbf{s}_1, \mathbf{s}_3 - \mathbf{s}_1, \dots, \mathbf{s}_n - \mathbf{s}_1]^T.$$

If each  $\mathbf{s}_i$  contains  $F$  ASM points, the final vector  $\mathbf{m}$  has dimensionality  $2F(n-1)$ .

Given a set of example training motions  $\mathbf{m}_i$  for a given disease and viewpoint, our linear model for a new motion  $\mathbf{m}$  is  $\mathbf{m} = \sum_i c_i \mathbf{m}_i$ . Applying PCA to the training set  $\mathbf{m}_i$  yields a set of eigenmotions  $\mathbf{e}_i^m$  and mean motion  $\bar{\mathbf{m}}$ , giving us a

lower dimensional representation

$$\mathbf{m} = M\mathbf{c} + \bar{\mathbf{m}} \quad M = \begin{bmatrix} | & | & & | \\ \mathbf{e}_1^m & \mathbf{e}_2^m & \dots & \mathbf{e}_r^m \\ | & | & & | \end{bmatrix},$$

where  $r$  eigenmotions are retained. The  $r$ -dimensional weight vector  $\mathbf{c}$  is our representation of object motion. Note that this representation is similar to Active Appearance Motion Models [1]. While their emphasis is to use the eigenmotions to assist/constrain tracking, we use motion eigenanalysis as a feature for disease classification.

To collect training motions, the manual labeling effort is significant since we must annotate entire sequences. We developed a semiautomatic approach that leverages our ASM tracker. We ran our tracker on a set of training sequences, evaluating the returned tracks and making manual corrections as necessary. After vectorizing the motion vectors and applying PCA, we can obtain motion models for our disease classes. Fig. 5 shows two example motion models from our experimental results in the next section; notice much greater motion for the normal class as compared to hypokinesia.

A new sequence  $\mathbf{m}$  can be analyzed to evaluate the degree of fit to our linear motion model. First, we project  $\mathbf{m}$  to find  $\mathbf{c}$

$$\mathbf{c} = M^T [\mathbf{m} - \bar{\mathbf{m}}]$$

and then the motion fit is a combination of the Mahalanobis distance and motion reconstruction

$$\text{m-fit}(\mathbf{c}) = \mathbf{c}^T \Sigma_{\text{mot}}^{-1} \mathbf{c} + 2R_{\text{mot}}^2 / \lambda_{\text{mot}}^{r+1} \quad (5)$$

where  $R_{\text{mot}} = \|\mathbf{m} - MM^T \mathbf{m}\|$ ,  $\lambda_{\text{mot}}^{r+1}$  is the  $(r+1)^{\text{th}}$  motion eigenvalue, and  $\Sigma_{\text{mot}}$  is a diagonal  $(r \times r)$  matrix with motion PCA eigenvalues. The Mahalanobis term is a weighted distance from the PCA projection of  $\mathbf{m}$  to the mean motion  $\bar{\mathbf{m}}$  (but within the PCA model space). The reconstruction term measures the distance from  $\mathbf{m}$  to the PCA projection of  $\mathbf{m}$ , and it tells how well the model explains the data.

#### D. Disease recognition algorithm

Putting it all together, Fig. 4 summarizes our algorithm for recognizing cardiac disease. We assume that all sequences used for training as well as testing represent single heart cycles, all synchronized at the peak of the R wave. The input to the algorithm are a set of sequences  $\text{Seq}_{i_1}, \text{Seq}_{i_2}, \dots$  from chosen viewpoints of a query patient with unknown cardiac disease, and a set of trained models per disease. When matching the query against a given disease  $\mathbf{D}_i$ , we first find a set of viewpoints  $K$  that they have in common. For each common viewpoint  $k \in K$ , we compute the fit between the input sequence  $\text{Seq}_k$  and models  $\text{ASM}_k$  and  $\text{Mot}_k$

$$\text{ViewFit}(i, k) = \frac{1}{F} \sum_{f=1}^F \text{fit}(\mathbf{a}_f, \mathbf{b}_f, \Gamma_f) + \text{m-fit}(\mathbf{c}). \quad (6)$$

The first term, from equation (4), averages the appearance fit over the  $F$  frames in  $\text{Seq}_k$ . The second term, from equation (5), adds in the motion fit from the entire sequence. Finally,

### Cardiac Disease Recognition Algorithm

#### Input:

Patient  $P$ : input sequences labeled by viewpoint  
 $P: \text{Seq}_{i_1}, \text{Seq}_{i_2}, \dots$  for viewpoints  $i_1, i_2, \dots$   
 $N$  disease models  $\mathbf{D}_i$ : ASM and motion models for  
 $\mathbf{D}_i$ 's diagnostic views  $j_1, j_2, \dots$   
 $\mathbf{D}_i: \text{ASM}_{j_1}, \text{Mot}_{j_1}, \text{ASM}_{j_2}, \text{Mot}_{j_2}, \dots$

#### Algo:

**for** disease  $i$  from 1 to  $N$  do  
 $K = \{i_1, i_2, \dots\} \cap \{j_1, j_2, \dots\}$   
**for** all views  $k \in K$  do  
1. Fit  $\text{ASM}_k$  to frame 1 of  $\text{Seq}_k$   
2. Track feature points in remainder of  $\text{Seq}_k$   
 $\rightarrow$  generates  $(\mathbf{a}_f, \mathbf{b}_f, \Gamma_f)$  over  $\text{Seq}_k$ ,  
 $1 \leq f \leq F, F = |\text{Seq}_k|$   
3. Project tracks onto  $\text{Mot}_k$   
 $\rightarrow$  generates  $\mathbf{c}$  motion vector  
4. Compute view fit  
 $\text{ViewFit}(i, k) = \frac{1}{F} \sum_{f=1}^F \text{fit}(\mathbf{a}_f, \mathbf{b}_f, \Gamma_f) + \text{m-fit}(\mathbf{c})$   
 $\text{DiseaseFit}(i) = \frac{1}{|K|} \sum_{k \in K} \text{ViewFit}(i, k)$   
**return**  $\underset{i}{\text{argmin}} (\text{DiseaseFit}(i))$

Fig. 4. Our algorithm for cardiac disease recognition.

the disease fitting score is the average over all  $k \in K$ , and the disease with minimum fitting score is reported.

The terms in our fitting score allow us to integrate information from a variety of sources. Spatial shape and textural appearance match information is incorporated into the ASM  $\mathbf{a}$  and  $\mathbf{b}$  terms in  $\text{fit}(\mathbf{a}, \mathbf{b}, \Gamma)$ . Motion similarity to the disease motion model is captured by the  $\mathbf{c}$  term in  $\text{m-fit}(\mathbf{c})$ . Finally, information from diagnostically important cardiac viewpoints are fused in the view averaging step to compute  $\text{DiseaseFit}(i)$ .

To represent the null hypothesis that the patient has no disease, we always add a ‘‘normal’’ class to model the non-disease state. The diagnostic viewpoints for the normal class are taken as the union of viewpoints from all the disease classes. This will always allow for the normal class to be compared against any of the disease classes when analyzing a new patient.

## IV. RESULTS

We now present experimental results on a 2 class problem: hypokinesia ( $\mathbf{D}_0$ ) versus normal patients ( $\mathbf{D}_1$ ). Hypokinesia is a cardiac condition where the heart suffers from reduced motion, so we expect the motion models to be quite different between the two classes. Fig. 5 shows the mean motion of normal patients versus Hypokinesia patients. It can be seen from this figure that the mean motion of normal patients is much greater than the hypokinetic mean. Thus, we expect the motion related term in the fitting metric to differentiate hypokinesia from normal patients.

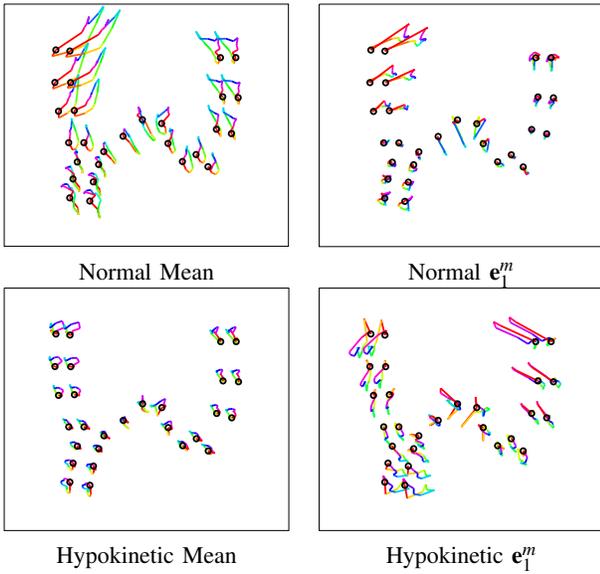


Fig. 5. Mean motion and first eigenmotion for normal patients and hypokinesia patients. Time position in the cycle is encoded using line segment hue. A large difference in mean motion intensity is immediately evident between normal and hypokinesia patients. (Best viewed in color.)

The diagnostic viewpoint we use for hypokinesia is Apical 4 Chamber (A4C). We built A4C models of hypokinesia and normal patients, and we matched these models against the A4C sequences of new patients with unknown disease labels.

Our data set came from hospitals in India with cardiologists recording a complete echo exam in continuous video. From these videos of the complete workflow, we manually extracted individual A4C cardiac cycles from both normal and hypokinesia patients. The echo “workflow” video was captured at 320x240 and 25 Hz, and a ECG trace waveform at the bottom allowed us to synchronize extracted cycles with the ECG R peak. Each video sequence was about 5 minutes long per patient and depicted several heart cycles. Our current collection has over 200 echo videos or 200x5x30x60 frames. The physicians in our project assisted us in their interpretation, so that the ground truth labels for the cardiac cycles could be obtained. Due to the effort in manual labeling for model training, we report our results on a data set with 16 hypokinesia patients and 5 normal patients.

The testing results for our disease recognition algorithm are given in Table I. During testing, because of our data set size, we employed a leave-one-out methodology: when testing a sequence from patient X, the entire data set minus X is used for training. From Table I, we notice that the recognition rate for normal class is lower than that of Hypokinesia class, which lowers the average recognition rate. Increasing the number of patients in the normal class will lead to better recognition rates.

While this hypokinesia vs. normal experiment emphasizes the motion term in equation (6), the shape appearance term in (6) would help distinguish diseases characterized by shape differences, such as the enlarged heart chambers in dilated cardiomyopathy.

	All patients (N = 21)	Hypokinesia $D_0$ (N = 16)	Normal $D_1$ (N = 5)	Disease Ave $(D_0 + D_1)/2.0$
Recog Rate	85.7%	93.7%	60.0%	76.8%

TABLE I  
DISEASE RECOGNITION RATE FOR HYPOKINESIA VS. NORMAL.

## V. CONCLUSION

In this paper, we have proposed a statistical appearance and motion modeling approach for directly detecting cardiac disease in echocardiographic sequences. We use Active Shape Models to represent heart shape and texture appearance in individual frames. To represent motion, we construct disease-specific PCA “eigenmotion” models based on the tracked motion of the ASM shape feature points. Disease recognition is posed as a model fitting problem where we fit the ASM and motion models trained from each disease to each new patient’s data and report the disease with best average fit over viewpoints. Overall, our framework allows us to integrate measures of shape, textural appearance, and motion, and it provides a means to fuse evidence across different echo viewpoints. Our initial experiments have shown the effectiveness of this approach in discriminating between Hypokinesia patients and normal patients. Future work will involve more extensive experiments on a large number of disease classes and viewpoints.

## REFERENCES

- [1] J. Bosch, S. Mitchell, B. Lelieveldt, F. Nijland, O. Kamp, M. Sonka, and J. Reiber. Automatic segmentation of echocardiographic sequences by active appearance motion models. *IEEE Transactions on Medical Imaging*, 21(11):1374–1383, November 2002.
- [2] T. Cootes and C. Taylor. Using grey-level models to improve active shape model search. In *International Conference on Pattern Recognition*, volume 1, pages 63–67, 1994.
- [3] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models—their training and application. *Comput. Vis. Image Underst.*, 61(1):38–59, 1995.
- [4] I. Dydenko, F. Jamal, O. Bernard, J. D’hooge, I. E. Magnin, and D. Friboulet. A level set framework with a shape and motion prior for segmentation and region tracking in echocardiography. *Medical Image Analysis*, 10(2):162–177, 2006.
- [5] A. Montillo, D. N. Metaxas, and L. Axel. Automated model-based segmentation of the left and right ventricles in tagged cardiac MRI. In *MICCAI*, pages 507–515, 2003.
- [6] G. Naylor, N. Firmin, and D. Longmore. Blood flow imaging by cine magnetic resonance. *J. Comp. Assist. Tomog.*, 10:715–722, 1986.
- [7] X. Papademetris, A. J. Sinusas, D. P. Dione, and J. Duncan. 3d cardiac deformation from ultrasound images. In *MICCAI*, pages 420–429, 1999.
- [8] N. J. Pelc. Myocardial motion analysis with phase contrast cine MRI. In *Proceedings of the 10th Annual SMRM*, 1991.
- [9] T. Syeda-Mahmood, F. Wang, D. Beymer, M. London, and R. Reddy. Characterizing spatio-temporal patterns for disease discrimination in cardiac echo videos. In *MICCAI*, pages 261–269, 2007.
- [10] T. Syeda-Mahmood and J. Yang. Characterizing normal and abnormal cardiac echo motion patterns. In *Computers in Cardiology*, pages 725–728, 2006.
- [11] A. Wong, P. Shi, H. Liu, and A. Sinusas. Joint analysis of heart geometry and kinematics with spatiotemporal active region model. In *EMBC*, pages 762–765, 2003.
- [12] E. Zerhouni, D. Parish, W. Rogers, A. Yang, and E. Shapiro. Human heart: Tagging with MR imaging – a method for noninvasive assessment of myocardial motion. *Radiology*, 169(1):59–63, 1988.