

Automatic Valve Segmentation in Cardiac Ultrasound Time Series Data

Yonatan Dukler¹, Yurun Ge², Yizhou Qian¹, Shintaro Yamamoto³, Baichuan Yuan¹, Long Zhao⁴, Andrea L. Bertozzi¹, Blake Hunter⁵, Rafael Llerena⁶, and Jesse T. Yen⁶

¹University of California, Los Angeles, USA

²Shanghai Jiao Tong University, China

³Waseda University, Tokyo, Japan

⁴University of California, Berkeley, USA

⁵Claremont McKenna College, Claremont, USA

⁶University of Southern California, Los Angeles, USA

ABSTRACT

We consider the problem of automatically tracking the mitral valve in cardiac ultrasound time series and present an unsupervised method for decomposing and segmenting the mitral valve from noisy ultrasound videos. To do so we propose a Robust Nonnegative Matrix Factorization (RNMF) method that naturally decomposes the time series into three separate parts, highlighting the cardiac cycle, mitral valve, and ultrasound noise. The low rank component of RNMF captures the simple motions of the cardiac cycle effectively aside from the sporadic motion of the mitral valve tissue that is captured innately in our RNMF sparse signal term. Using the RNMF representation, we introduce a simple valve object detection algorithm. Our method performs especially well in noisy time series when existing methods fail, differentiating general noise from the subtle and complex motions of the mitral valve. The valve is then segmented using simple thresholding and diffusion. The method presented is highly robust to low quality ultrasound video, and does not require manual preprocessing, prior labeling, or any training data.

Keywords: Robust Nonnegative Matrix Factorization, Cardiac Ultrasound, Valve Detection, Motion segmentation

1. INTRODUCTION

Echocardiography is a type of ultrasound test that uses high-pitched sound waves to produce an image of the heart and provides valuable information for the assessment of heart diseases and abnormalities. Ultrasound images from echocardiography are commonly used^{1,2} to determine the left ventricular ejection fraction (LVEF). In the medical field today, highly trained specialists manually process these images by segmenting features of interest. Due to the considerations of time, cost, and distance from specialists, the significance of automatic segmentation of features in ultrasound images cannot be overstated.

The mitral valve, a dual-flap valve in the heart that lies between the left atrium and the left ventricle, plays an important role in providing medical information about a plethora of valvular diseases³ and in potentially assessing the left ventricular systolic function.⁴ Currently, the delineation of the mitral valve is a time-consuming process manually done by a medical practitioner. Therefore, an automatic object tracking algorithm for the mitral valve is of great interest.

Many methods for mitral leaflet segmentation tracks the valve are based on active contours.⁵ When such methods segment the valve, they treat each frame separately resulting in tedious user involvement to initialize each frame or to supervise the partial results to assure accurate segmentation results.⁶ More recently, low-rank sparse decomposition techniques have been used for object tracking algorithms. Such algorithms have

Corresponding author: Baichuan Yuan
Email address: byuan@math.ucla.edu

been designed to automatically track the mitral valve in ultrasound time series using a low-rank approximation, utilizing the motion complexity of the mitral valve.⁷ Such low rank approximation methods, however, are often misled by the particle noise presents in ultrasound videos and must improve in robustness. Such is the innovative method presented in Ref. 7. for the unsupervised tracking of the mitral valve based on low-rank representation. However, due to the high noise-level in echocardiography videos, the method is not robust enough to withstand the speckle disturbance that is present in a wide range of ultrasound image sequences.

In this paper, we create an automatic algorithm to segment the mitral valve with low computational cost using Robust Nonnegative Matrix Factorization (RNMF). In addition we use our RNMF framework to create an improved mitral valve detection algorithm similar to that presented in Ref. 7. Our method does not rely on users to manually adjust parameters tailored for each specific patient. On the contrary, only a generic window size is required for a large collection of videos with a similar view of the heart. Instead of regarding all frames from a single video as unrelated static images our method exploits the continuity of the video frames as a time series. The key idea used in separating the mitral valve from the rest of the heart is using its motion complexity that allows for separation of the valve’s non-linear motion from the generally low-rank ultrasound video. The periodic low-rank expansion and contraction of the heart can be naturally captured by a low rank term in RNMF whereas the mitral valve is captured as the high dimensional sparse term. Comparing to other dimensionality reduction algorithms, RNMF provides highly meaningful end-members and corresponding coefficients in the realm of echocardiography.

This paper extend on previous work⁸ done on low rank cardiac ultrasound that illustrated the advantage of a rank-2 NMF method for the approximation of the cardiac cycle. NMF is a linear dimensionality reduction method⁹ that is superior in non-negative datasets and allows for great interpretation of results. In the context of cardiac ultrasound, rank-2 NMF approximates the heart cycle form cardiac ultrasound videos by describing each frame as a non-negative mixture between two end-member frames. Each end-member generalizes the end systole (fully contracted) or end diastole (fully expanded) stages of the heart cycle. In this paper we continue using such framework but recognize that certain tissues of the heart, namely the mitral valve, are not captured effectively by such low rank approximations and therefore improve the framework by using Robust NMF as opposed to classic NMF. We illustrate that the new framework preserves the desired properties found in NMF. In addition, the new robust framework is able to detect the mitral valve in noisy environments and separate it from the low rank approximation inspired by classical NMF.

Our mitral valve segmentation algorithm has the immediate application of calculating the mitral annular displacement. By selecting three points of interest according to the segmentation result and the original video, we can approximate the mitral annular displacement. Recent research⁴ on the relationship between the left ventricular ejection fraction and mitral annular displacement demonstrates a strong curvilinear relationship by statistical fitting of data from patients’ cardiac cycle in the apical 4- and 2- chamber views. With this correlation at hand, we present an automatic algorithm to estimate ejection fraction via our segmentation of the mitral valve.

2. MATERIAL AND METHODS

2.1 Source of Data

A total of 99 echocardiogram samples in the apical 4- and 2- chamber view were collected randomly from a database at the Keck Medical Center at the University of Southern California, and properly de-identified for the analysis. All videos and relevant diagnosis data were obtained following a Health Sciences Review Board (HSIRB) approved protocol HS-15-00258. HSIRB reviews all human subject research protocols in accordance with federal regulation, state law, and university policy. The ultrasound images were originally obtained from the GE Vivid E9 ultrasound system under the same settings prior to be transferred to the database. The samples include a diverse group of abnormal conditions and a few normal studies to review the anatomical structures pertinent to each condition and be subject of the analysis. The echocardiogram were obtained from pathologic conditions that have different effects in the anatomical structures to be analyzed. These conditions included ischemic cardiomyopathy, hypertrophic cardiomyopathy, heart failure, aortic valve stenosis, myocardial infarction, pulmonary hypertension, systemic hypertension, myocardial toxicity, tamponade and congenital heart disease.

Based on the quality of videos, we classify the selected 93 videos from our dataset into three categories: high quality, medium quality, and low quality. There are 16 videos with high quality, where the presence of the mitral valve is clear and where the noise level is low. There are 23 videos with medium quality, where the mitral valve is visible but notable noise is present. The rest 54 videos are of low quality, where the valve is vague and the noise level is high. For the purpose of valve segmentation, we disregard 6 videos in our dataset where the mitral valve is completely absent.

2.2 Robust Nonnegative Matrix Factorization (RNMF)

RNMF is an extension of the dimensionality reduction method NMF. It approximates a nonnegative large matrix as the sum of the product of two nonnegative low-rank matrices and a sparse noise matrix. Rank-2 NMF is used in Ref. 8 to model the muscle movement as a linear combination of two end-members, which correspond to end-systole and end-diastole stages of the cardiac ultrasound. In RNMF an important underlying assumption is that the data matrix \mathbf{X} is generally of low rank k aside from a sparse component \mathbf{S} of assumed higher rank. Via RNMF the nonnegative data matrix $\mathbf{X} \in \mathbb{R}_+^{m \times n}$ is factored into

$$\mathbf{X} \approx \mathbf{W}\mathbf{H} + \mathbf{S} \quad (1)$$

with $\mathbf{W} \in \mathbb{R}_+^{m \times k}$, $\mathbf{H} \in \mathbb{R}_+^{k \times n}$, $\mathbf{S} \in \mathbb{R}_+^{m \times n}$ and \mathbf{S} being sparse. Here, $\mathbf{W}\mathbf{H}$ is a matrix of rank of at most k , being the product of two matrices of rank k , and \mathbf{S} is the sparse signal. To achieve the desired decomposition, one needs to minimize the energy function

$$f(\mathbf{W}, \mathbf{H}, \mathbf{S}) = \|\mathbf{X} - \mathbf{W}\mathbf{H} - \mathbf{S}\|_F^2 + \lambda \|\mathbf{S}\|_1. \quad (2)$$

Intuitively speaking, the Frobenius norm $\|\cdot\|_F$ enforces the proximity of \mathbf{X} to $\mathbf{W}\mathbf{H} + \mathbf{S}$, and the $\|\cdot\|_1$ norm on \mathbf{S} promotes the sparsity of the matrix \mathbf{S} , where the sparsity parameter λ controls the extent to which sparsity is enforced.

We intend on applying rank-2 RNMF to decompose the cardiac ultrasound time series as a low-rank approximation of the heart muscles and a sparse high dimensional term capturing the mitral valve. Initially the ultrasound cardiac video is represented by a tensor $\mathbf{T} \in \mathbb{R}_+^{r \times c \times l}$ consisting of a sequence of images $\mathbf{T}(\cdot, \cdot, i) \in \mathbb{R}_+^{r \times c}$. To apply RNMF one needs to be convert \mathbf{T} into a matrix. We reshape each image $\mathbf{T}(\cdot, \cdot, i) \in \mathbb{R}_+^{r \times c}$ into a column vector $x_i \in \mathbb{R}_+^{(r \cdot c) \times 1}$. All such column vectors are then concatenated together to produce an RNMF-ready nonnegative data matrix \mathbf{X} . We follow an iterative thresholding approach¹⁰ to solve our RNMF framework.

Algorithm 1 SOLVING ROBUST NONNEGATIVE MATRIX FACTORIZATION USING ITERATIVE THRESHOLDING

Input: Nonnegative matrix $X \in \mathbb{R}_+^{n \times m}$, rank $k = 2$, sparsity coefficient λ , tolerance ϵ , max iterations N

Output: Nonnegative matrices $W \in \mathbb{R}_+^{n \times k}$, $H \in \mathbb{R}_+^{k \times m}$, $S \in \mathbb{R}_+^{n \times m}$

1: $i \leftarrow 0$, initialize W^0 and H^0 randomly at uniform or using NMF.

2: **while** $i \leq N$ and $|f(W_{i-1}, H_{i-1}) - f(W_i, H_i)| > \epsilon$ **do**

3: $S^{i+1} \leftarrow X - W^i H^i$

4: $S_{ij} \leftarrow \begin{cases} S_{ij} - \frac{\lambda}{2} & S_{ij} > \frac{\lambda}{2} \\ 0 & \text{otherwise} \end{cases}$

5: $W^{i+1} \leftarrow W^i \circ \frac{(X - S^{i+1})_+ H^T}{W H H^T}$

6: $H^{i+1} \leftarrow H^i \circ \frac{W^T (X - S^{i+1})_+}{W^T W H}$

7: $W^{i+1} \leftarrow \frac{W^{i+1}}{\|W^{i+1}\|_F}$

8: $H^{i+1} \leftarrow \frac{H^{i+1}}{\|H^{i+1}\|_F}$

9: $i \leftarrow i + 1$

10: **end while**

11: Output W_i, H_i, S_i

In the above and throughout the paper the truncation operator $(x)_+$ refers to $(x)_+ := \begin{cases} x & x > 0 \\ 0 & x < 0 \end{cases}$ and \circ refers to the entry-wise matrix product.

2.3 Automatic Valve Detection

A direct application of RNMF is the object detection – that is to find a rectangular region (window) bounding the desired object – of the mitral valve. In Ref. 7 the desired window is selected as the window of maximum singular values tail sum, indicating of complex non-linear motion. Inspired by this idea, we propose an RNMF-based approach in Algorithm 2, which selects the desired window as the window with the largest Frobenius norm on the matrix \mathbf{S} in RNMF. For the purpose of completeness, we also include the SVD-based window detection algorithm used in 7. In the following algorithm, RNMF method refers to our newly designed method, and SVD method refers to the existing approach.

Algorithm 2 Automatic Window Detection

Input: Data matrix $M = (I_1, \dots, I_n)$, the set of all possible rectangular windows $\{W_j\}$, and singular value threshold $K > 0$ (only for SVD method)

Output: Detected window W_{j^*}

- 1: **for** each spatial window W_j **do**
 - 2: let $M_j = M \circ W_j$
 - 3: for SVD method: compute the residual singular values $\sigma_1, \dots, \sigma_n$ for M_j , and
 $\epsilon_j = \sum_{i=K+1}^n \sigma_i^2$
 - 4: for RNMF method: obtain the approximation $M_j \approx W_j H_j + S_j$, and
 $\epsilon_j = \|S_j\|_F$
 - 5: **end for**
 - 6: let $j^* = \arg \max_j \epsilon_j$
-

In Algorithm 2, each window W_j is a rectangular region with fixed size and position in the video. The size of the spatial window is fixed for the entire data-set and is of the approximate size of the mitral valve. The positions of the spatial windows start from the top left corner of the video and end at the bottom right corner, with specified step sizes. For each window, either SVD or RNMF is performed. We then select the window that maximizes nonlinearity measured either by the Frobenius norm of the sparse matrix (\mathbf{S}) or by the sum of the tail singular values starting from σ_{K+1} . In both options the calculated value is indicative of the dimensionality of the window. The fast changing mitral valve region exhibits great dimensional complexity and is therefore captured by such measures.

2.4 Ultrasound Time Series Data Decomposition

As mentioned before, the low rank approximation \mathbf{WH} from RNMF highlights the myocardium movement while the sparse term \mathbf{S} from RNMF illustrates the motion of the mitral valve. Following this key idea, we present an automatic segmentation method in Algorithm 3, which decomposes the original video into three components: (1) myocardium, (2) mitral valve and (3) residual noise. In this algorithm, a thresholding function T_p and an anisotropic diffusion operation¹¹ D_r are used. The thresholding function T_p ranks the pixels in terms of brightness, and returns a logical matrix consisting of the top $p\%$ pixels. The anisotropic diffusion operator D_r connects small disjoint components and smooths the boundary of segmentation.

Algorithm 3 Mitral Valve Decomposition

Input: image sequence $M = (I_1, \dots, I_n)$, thresholding function T_P , diffusion function D , percentage parameters P_1, P_2

Output: myocardium W_2H_2 , mitral valve R and noise $M - W_2H_2 - R$

- 1: Apply RNMF to obtain $M \approx W_1H_1 + S_1$
 - 2: Apply the Window Detection Algorithm 2 to detect the region W_{j^*}
 - 3: Calculate Image sequence without valve $M' = M \circ (1 - S')$, where $S' = D_r(W_{j^*} \circ T_{P_1}(S_1))$.
 - 4: Apply RNMF again to obtain muscle motion $M' \approx W_2H_2 + S_2$
 - 5: Find valve $R = D_r(W_{j^*} \circ T_{P_2}(M - W_2H_2))$
-

By applying the thresholding function T_p to the resulting sparse matrix S , we are able to obtain $T_{p_1}(S)$ that highlights the pixels corresponding to the most complex and nonlinear motion of the mitral valve, as well as some noise from the myocardium. We then remove the noise by restricting our result to the window obtained by the window detection Algorithm 2. Furthermore, using anisotropic diffusion on the data allows us to ensure the connectivity of our valve segmentation and accurate detection of the boundary of the tissue. We then subtract the temporary segmentation result S' from the original data to highlight the muscle motion. To provide a low rank approximation of the muscle motion we apply RNMF one more time to obtain the low-rank term W_2H_2 . The resultant residue S_2 then fully highlights the mitral valve including tissue neglected by the first sparse term S_1 . Therefore we obtain our final mitral valve segmentation R by applying thresholding and anisotropic diffusion on S_2 again in order. Finally, we obtain the residual noise term simply by subtracting $W_2H_2 + R$ from original data.

2.5 LVEF Estimation from Mitral Annular Displacement

Automatic estimation of LVEF is of great interest in medical image processing. One important application of the valve-muscle decomposition in Algorithm 3 is in calculating the mitral annular displacement (MD). According to recent research,⁴ LVEF is highly correlated with the MD between the two cardiac end-stages. To compute MD, one needs to trace three key points consisting of the apical myocardium and the hinge points of the mitral valve leaflets with the septal and lateral aspects of mitral annulus.

In Algorithm 4, we detect the apical myocardium by image morphology and a moving-window threshold function. For the two hinge points, we utilize our segmentation results to find the leftmost and rightmost end points of the mitral valve. Then we track the movement of these points by optical flow.¹² The displacement of the midpoint between the two hinge points towards the apex point is recorded. To account for different scales of the left ventricle, the displacement of the two hinge points is normalized by the chamber length during cardiac end-diastole. Figure 1 displays the locations of the three points of interest.

Algorithm 4 Apical Myocardium and Hinge Points Detection

Input: image sequence $M = (I_1, \dots, I_n)$, valve segmentation result $V = (V_1, \dots, V_n)$

- 1: **for** each j **do**
 - 2: Binarize I_j as T_j using window shift thresholding
 - 3: Apply morphological closure of T_j with disc of diameter 5
 - 4: **{Apical Myocardium tracking:}**
 - 5: Find the third row with non-zero elements from top of T_j , Record row number as p
 - 6: Find the midpoint of the non-zero elements in row p , record the column number of this point as topcol_j .
 - 7: Select the first zero element in col topcol_j starting from row p , record the row number of this element as toprow_j
 - 8: **{hinge points tracking:}**
 - 9: Find the third column with non-zero elements in V_j from the left side, Record column as q_1
 - 10: Find the midpoint (leftrow $_j$, leftcol $_j$) of the non-zero elements in column q_1 of V_j
 - 11: Find the third column with non-zero elements in V_j from the right side, Record column as q_2
 - 12: Find the midpoint (rightrow $_j$, rightcol $_j$) of the non-zero elements in column q_2 of V_j
 - 13: **end for**
- Output:** Output the coordinate of apical myocardium ($\text{toprow}_j, \text{topcol}_j$), left hinge point (leftrow $_j$, leftcol $_j$) and right hinge point (rightrow $_j$, rightcol $_j$) for each j
-

The third pixel is chosen in Algorithm 4 to avoid possible outliers and does not pose other significant discrepancies.




Figure 1: The top left image illustrates the locations of the apical myocardium and the hinge points of the mitral valve leaflets with the septal and lateral aspects of mitral annulus. The valve segmentation result is used to determine the locations of the hinge points. The picture on the top right represents how mitral annulus displacement (MD) is measured between the two end-member phases. The bottom image demonstrates how the chamber length is calculated during cardiac end-diastole.

3. RESULT

3.1 Robust Nonnegative Matrix Factorization Properties

In NMF the residue term $\mathbf{X} - \mathbf{WH}$ was shown to highlight the movement of the mitral valve since its complex, nonlinear motion cannot be captured by a low-rank decomposition. Inspired by this observation, this paper introduces a sparse term \mathbf{S} in RNMF which captures complex motions in the data matrix \mathbf{X} . Unlike the residue term $\mathbf{X} - \mathbf{WH}$ in NMF, which has negative entries, the components of \mathbf{S} are non-negative and only capture physical objects. In Figure 2, we illustrate the great extent of noise previously present in the residue term $\mathbf{X} - \mathbf{WH}$ which is eliminated in the sparse term \mathbf{S} . Nonetheless, the mitral valve structure and texture of the valve remain.




Figure 2: We present two frames to illustrate the sparse term \mathbf{S} in RNMF side by side with the NMF residual $\mathbf{X} - \mathbf{WH}$ in absolute value. \mathbf{S} consists of 6% of the total pixels, while $\mathbf{X} - \mathbf{WH}$ consist of 40% of all the pixels in the video.

An important property of rank-2 NMF in the realm of echocardiography is its effective capture of the cardiac cycle. That is, the coefficient matrix \mathbf{H} describes the nonnegative mixture of the video frames and the heart cycle and can lead to automatic detection of the frames of end-systole and end-diastole.⁸ We illustrate that this property is preserved when one switches from NMF to its robust counterpart, RNMF. In an NMF decomposition each column vector of a matrix \mathbf{W} is said to be an end-member, and the corresponding row vector of the matrix \mathbf{H} represents the nonnegative linear coefficients for the corresponding end-members across the given time series. In rank-2 RNMF, the matrices \mathbf{W} and \mathbf{H} produce the same desired interpretation as in rank-2 NMF. In particular, when we apply rank-2 (R)NMF to ultrasound cardiac time series, the two end-members of (R)NMF are the frames corresponding to end-systole and end-diastole.

To understand the periodic pattern of the coefficients in \mathbf{H} , we plot the normalized coefficients over time in both RNMF and NMF in Figure 3. The blue and orange curves represent the coefficients for the first and second end-members of \mathbf{W} respectively. To identify the frames corresponding to the end-systole and end-diastole, we select the frames at which the blue or the orange curves attain their maximum values. That is, we select the two frames in the time series that most resemble one of the end-members in \mathbf{W} . Despite some small discrepancies in the two graphs, the results of automatic detection of end-systole and end-diastole are not affected. The meaningful cyclic properties of rank-2 NMF are therefore preserved in rank-2 RNMF.

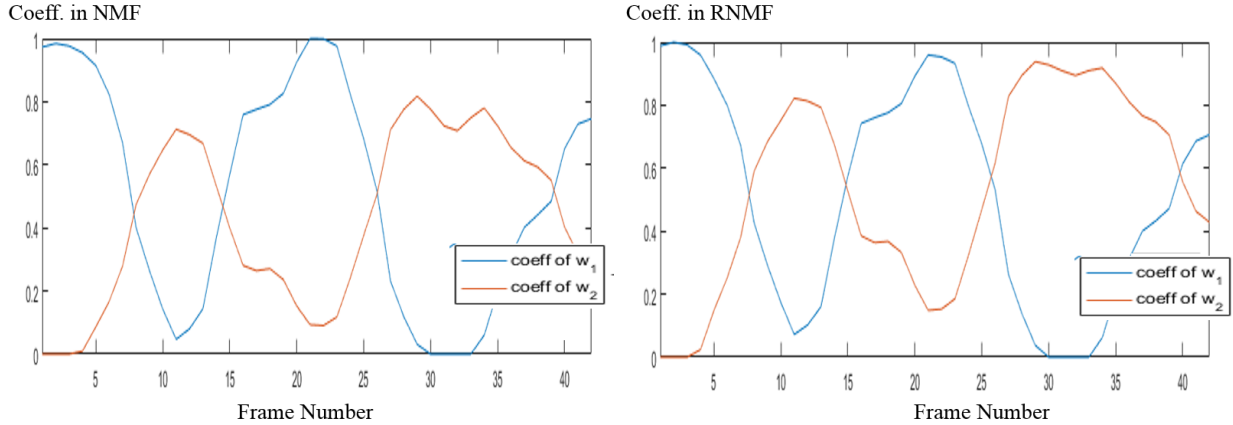


Figure 3: Plots illustrating the non-negative linear coefficients present in \mathbf{H} . The periodic pattern of the coefficients is preserved in RNMF as in NMF.

3.2 Window Detection

In the following tables, we present the window detection results from Algorithm 2 based on the three video quality categories we discussed above. To help understand the performance of RNMF-based window detection algorithm, we compare its results with the SVD-based algorithm used in Ref. 7.

We observe that the SVD-based algorithm produces accurate results when the videos are of high quality. However, its accuracy drops significantly when processing videos with medium or low quality. On the other hand, RNMF performs well under high-quality videos and maintains its good performance in the cases of medium-quality videos. Despite an accuracy drop for low-quality videos, RNMF still outperforms the SVD-based algorithm significantly.

Table 1: Quantitative evaluation for mitral valve detection between methods. The window method with a fixed size across all cases is successful if more than 65% of the window is contained in the manually-generated bounding box for the valve.

High Quality Videos	number of successful cases	Total cases	Success rate (%)
RNMF on window	15	16	94
SVD on window	15	16	94
Medium Quality Videos	number of successful cases	Total cases	Success rate (%)
RNMF on window	21	23	91
SVD on window	17	23	74
Low Quality Videos	number of successful cases	Total cases	Success rate (%)
RNMF on window	30	54	56
SVD on window	26	54	48

In particular, we should point out that the results of the SVD are likely affected by the high intensity of noise in the videos. In Figure 4, we demonstrate the resilience of the rank-2 RNMF technique to varying pixel intensity as compared to SVD. This might explain why the SVD-based algorithm experiences a considerable drop in accuracy with reduced video quality while RNMF maintains better results.

3.3 Mitral Valve Segmentation

In Algorithm 3, we decompose the original video into mitral valve, myocardium, and residual noise. The decomposition uses RNMF results and applies the threshold and diffusion operators. Figure 5 shows the video decomposition results from four cases. For the purpose of comparison, we manually trace the valve motion and



Figure 4: Resulting selected windows from the valve detection algorithm (Algorithm 2) of our proposed RNMF method (green) and SVD method (red) in Apical 4 chamber views.

create a ground truth (GT) segmentation for the mitral valve. In addition, the existing algorithm by Ref. 7 is also included for comparison. Both of the methods are automatic and do not require additional manual adjustments from the user.

Table 2: Four high-quality videos are chosen randomly from our dataset. In all of the three metrics, our method outperforms the current method⁷ with the only exception of precision in Case B.

Algorithm 3 (RNMF Based)	Precision (%)	Recall (%)	MAD (mm)
Case A	0.86 ± 0.07	0.63 ± 0.12	0.81 ± 0.21
Case B	0.80 ± 0.07	0.64 ± 0.12	0.69 ± 0.22
Case C	0.78 ± 0.12	0.66 ± 0.18	0.69 ± 0.23
Case D	0.75 ± 0.15	0.60 ± 0.18	0.94 ± 0.99
Overall	0.80 ± 0.10	0.63 ± 0.15	0.78 ± 0.41
Method in Ref. 7	Precision (%)	Recall (%)	MAD (mm)
Case A	0.77 ± 0.09	0.39 ± 0.07	1.49 ± 1.21
Case B	0.86 ± 0.12	0.35 ± 0.08	0.93 ± 0.29
Case C	0.42 ± 0.26	0.15 ± 0.10	3.82 ± 2.37
Case D	0.43 ± 0.25	0.31 ± 0.17	4.71 ± 3.64
Overall	0.62 ± 0.18	0.30 ± 0.11	2.74 ± 1.88

We compare both results with the ground truth in three metrics: precision, recall and mean absolute distance (MAD). The definitions are as follows:

$$\text{precision} = \frac{|R_{\text{alg}} \cap R_{\text{GT}}|}{|R_{\text{alg}}|}, \quad \text{recall} = \frac{|R_{\text{alg}} \cap R_{\text{GT}}|}{|R_{\text{GT}}|},$$

$$\text{MAD}(C_{\text{alg}}, C_{\text{GT}}) = \frac{1}{2} \int_0^1 [d(C_{\text{alg}}(s), C_{\text{GT}}) + d(C_{\text{GT}}(s), C_{\text{alg}})] ds.$$

Here, R_{alg} denotes the algorithm-generated mask for segmentation, R_{GT} stands for the manually-traced segmentation of the mitral valve, and $|S|$ is defined as the number of pixels inside the region S . $C(s)$ is a parametrization of contour C , namely the boundary of R , with constant velocity $|C'(s)|$, and $d(p, C) = \min_q \|p - C(q)\|_2$ is the minimum distance from a point p to a contour C . The higher the value in precision and recall, the closer the segmentation is to the ground truth. Meanwhile, the lower the value in MAD, the closer the contour of segmentation is to that of the ground truth. In Table 2, four high-quality videos are chosen randomly from our dataset. In all of the three metrics, our method outperforms the method in Ref. 7 except for the precision in Case B.




Figure 5: Video decomposition results from four random cases of our data. Three frames are selected for each video. The three columns from left to right represent original image, myocardium, and mitral valve segmentation respectively.

3.4 LVEF Estimation

We estimate the ejection fraction of patients using the mitral valve excursion as presented in Ref.13. The recent research suggests that there exists a curvilinear relationship between MD and LVEF ($R^2 = 0.642$). The following correlation is achieved by statistical fitting from 266 patients:⁴

$$LVEF = -4.256 + 8.263 * MD - 0.242 * MD^2. \quad (3)$$

Table 3: EF estimation by mitral valve excursion using equation (3).

	MD	EF	GT (EF)
Case 1	11.91	59.84	58
Case 2	18.18	65.98	63
Case 3	17.31	66.26	62
Case 4	10.97	57.26	65

4. FUTURE WORK

For future work, one could estimate the Ejection Fraction (EF) using the mitral valve excursion. Employing our segmentation result, we are also interested in developing a direct method for the EF using Simpson's biplane method: by segmenting the mitral valve from the cardiac video one is able to obtain a much simpler time series for the detection of the heart walls used in the calculation of EF. We are interested in expanding such detection in order to build algorithms for the detection of the boundaries of the left ventricle which will be outside of our mitral valve segmentation.

In addition, the current algorithm does not use the temporal information of the video sequence. To improve that, we could exploit the spatio-temporal connectivity of the valve tissue. It is also necessary to change the valve window size according to the scale of each video as we fix the valve window size for all scale videos. However, it is a challenging task because our algorithm measures the scale of each video based on the chamber length which is measured using two hinge points estimated from the valve segmentation result. Therefore, we need to measure the scale of each video by another algorithm that does not depend on the window detection algorithm.

ACKNOWLEDGMENTS

We thank Richard Koffler of Viderics, Inc. for useful comments. This work was supported by UCLA through the Physical Sciences Division including the Entrepreneurship and Innovation Fund and the Department of Mathematics. ALB, BH, YD and YQ were supported by NSF grant DMS-1417674 and ONR grant N00014-16-1-2119. SY and YG were supported by the Cross-disciplinary Scholars in Science and Technology (CSST) program at UCLA.

REFERENCES

- [1] Schiller, N. B., Shah, P. M., Crawford, M., DeMaria, A., Devereux, R., Feigenbaum, H., Gutgesell, H., Reichek, N., Sahn, D., Schnittger, I., et al., "Recommendations for quantitation of the left ventricle by two-dimensional echocardiography," *Journal of the American Society of Echocardiography* **2**(5), 358–367 (1989).
- [2] Pombo, J. F., Troy, B. L., and Russell, R. O., "Left ventricular volumes and ejection fraction by echocardiography," *Circulation* **43**(4), 480–490 (1971).
- [3] Nkomo, V. T., Gardin, J. M., Skelton, T. N., Gottdiener, J. S., Scott, C. G., and Enriquez-Sarano, M., "Burden of valvular heart diseases: a population-based study," *The Lancet* **368**(9540), 1005–1011 (2006).
- [4] Suzuki, K., Akashi, Y. J., Mizukoshi, K., Kou, S., Takai, M., Izumo, M., Hayashi, A., Ohtaki, E., Nobuoka, S., and Miyake, F., "Relationship between left ventricular ejection fraction and mitral annular displacement derived by speckle tracking echocardiography in patients with different heart diseases," *Journal of cardiology* **60**(1), 55–60 (2012).
- [5] Isard, M. and Blake, A., "Contour tracking by stochastic propagation of conditional density," in [*European conference on computer vision*], 343–356, Springer (1996).
- [6] Mikic, I., Krucinski, S., and Thomas, J. D., "Segmentation and tracking in echocardiographic sequences: Active contours guided by optical flow estimates," *IEEE transactions on medical imaging* **17**(2), 274–284 (1998).
- [7] Zhou, X., Yang, C., and Yu, W., "Automatic mitral leaflet tracking in echocardiography by outlier detection in the low-rank representation," in [*Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*], 972–979, IEEE (2012).
- [8] Yuana, B., Chitturib, S. R., Iyera, G., Lia, N., Xud, X., Zhanc, R., Llerenae, R., Yene, J. T., and Bertozzia, A. L., "Machine learning for cardiac ultrasound time series data," in [*SPIE Medical Imaging*], 101372D–101372D, International Society for Optics and Photonics (2017).
- [9] Lee, D. D. and Seung, H. S., "Learning the parts of objects by non-negative matrix factorization," *Nature* **401**(6755), 788–791 (1999).
- [10] Zhang, L., Chen, Z., Zheng, M., and He, X., "Robust non-negative matrix factorization," *Frontiers of Electrical and Electronic Engineering in China* **6**(2), 192–200 (2011).

- [11] Perona, P. and Malik, J., “Scale-space and edge detection using anisotropic diffusion,” *IEEE Transactions on pattern analysis and machine intelligence* **12**(7), 629–639 (1990).
- [12] Farneäck, G., “Two-frame motion estimation based on polynomial expansion,” *Image analysis* , 363–370 (2003).
- [13] Hu, K., Liu, D., Herrmann, S., Niemann, M., Gaudron, P. D., Voelker, W., Ertl, G., Bijnens, B., and Weidemann, F., “Clinical implication of mitral annular plane systolic excursion for patients with cardiovascular disease,” *European Heart Journal–Cardiovascular Imaging* **14**(3), 205–212 (2012).