

Directional Representations of 4D Echocardiography for Temporal Quantification of LV Volume

Elsa Angelini¹, Andrew Laine¹, Shin Takuma², Shunichi Homma²

¹ Department of Biomedical Engineering
Fu Foundation School of Engineering and Applied Science
Columbia University, New York, NY
laine@bme.columbia.edu

² Department of Medicine, Echocardiography Laboratories
College of Physicians and Surgeons, Columbia-Presbyterian Medical Center
Columbia University, New York, NY

Abstract. Real-time acquisition via four-dimensional (3D plus time) ultrasound obviates the need for slice registration and reconstruction, leaving segmentation as the only barrier to an automated, rapid, and clinically applicable calculation of accurate left ventricular cavity volumes and ejection fraction. Speckle noise corrupts ultrasound data by introducing sharp changes in an image intensity profile, while attenuation alters the intensity of equally significant cardiac structures, depending on orientation with respect to the position of the ultrasound beam. These properties suggest that measures based on phase information rather than intensity are appropriate for denoising and boundary (surface) detection. Our method relies on the expansion of temporal volume data on a family of basis functions called Brushlets. These basis functions decompose a signal into distinct patterns of oriented textures. Projected coefficients are associated with distinct “brush strokes” of a particular size (width) and orientation (direction). Brushlet decompositions are invariant to intensity (contrast range) but depend on the spatial frequency content of a signal. Preliminary results of this directional space-frequency analysis applied to both phantoms and clinical data are presented. The method will be used to clinically evaluate 4D data and to extract and quantify heart LV volumes.

1 Introduction

Cardiologists are excited about the opportunity for improved clinical and diagnostic performance via new 3D and 4D cardiac acquisition techniques, as they can view a heart at any angle with only one acquisition sequence. This represents a great potential cost savings for health care management as a patient would not need to be recalled in case a traditional echoplanar view did not reveal some crucial aspect of heart function. Also, cardiologists believe that a new visualization tool for displaying volumes in time would improve their diagnostic accuracy regarding tissue characterization and performance measures, such as cardiac output (CO) and ejection fraction (EF), by reducing the (inter-physician) variability of these measures.

Motivation: A diversity of modalities allows for the acquisition of dynamic sequences of cardiac volumes such as CT, Tagged MRI, SPECT and ultrasound. Echocardiography is the fastest, least expensive, and least invasive method for imaging the heart. The simplest and most useful clinical parameter used to assess cardiac function is ejection fraction (EF), calculated as the difference between end diastolic and end systolic left ventricular volumes. However, accurate calculations of ventricular volume from standard echocardiographic data are tedious and costly to employ clinically. This is because existing methods require time to digitize endocardial borders on a series of two-dimensional images, then register the image set and reconstruct each cavity volume.

Real-time acquisition via three-dimensional ultrasound obviates the need for slice registration and reconstruction, leaving segmentation as the remaining barrier to an automated, rapid, and therefore *clinically applicable calculation of accurate left ventricular cavity volumes and ejection fraction*.

Because it provides such a rich description of the temporal and spatial environment of any area of interest, three-dimensional ultrasound also offers the potential for *increased sensitivity* in detecting subtle wall motion abnormality indicative of ischemia (for example during an exercise stress test), compared to fast MRI techniques.

Existing Methods: State of the Art. Extraction of cardiac volume and quantitative analysis of wall deformation is of great interest to researchers in the field of medical imaging. Multigated radionuclide angiography (MUGA) and 2D echocardiography (2DE) are traditional screening techniques used in echocardiography. The main drawback of MUGA is that it requires the injection of a radiopharmaceutical agent. The 2D echocardiography technique is limited by its geometric resolution and by low signal-to-noise ratios (SNR) intrinsic to ultrasound images. Advanced invasive techniques such as Xray-CT and high cost nuclear modalities such as Tagged MRI and SPECT offer excellent resolution in space and time.

The majority of the volume extraction methods are based on prior models of the entire heart or of the left ventricle only. The parametrization of the model generally uses Finite Element models where the volume is constructed after deformation of the model following physics based constraints for equilibrium. Movement of the cardiac wall extracted from temporal data requires some parametrization of the model. Duncan et al. used contour shape descriptors in [1], Ayache; Cohen et al. [2] have used superquadratics. The nature of the constraints varies between models and can take a wide range of properties, including differential constraints [3], displacement and velocity constraints [4] as well as other constraints allowing for non-rigid movements.

Temporal Quantification: A dynamic Measure of CO. As a precursor for volume extraction, our data required some preprocessing to increase available SNR. Indeed it seems that any effort to build a 3D model and apply some deformation to isolate the LV volume from the raw data would be pointless because of the poor signal available for feature analysis. Following this statement, we employed multiscale denoising as a preprocessing step to volume extraction. The design principle of our denoising method relies on the expansion of temporal volume data on a family of basis

functions called Brushlets, introduced in 1997 by Meyer and Coifman [5]. These basis functions offer a decomposition of a signal into distinct patterns of oriented textures. In 2D, depending on the tiling of the L^2 domain chosen prior to analysis, the projected coefficients are associated with distinct “brush strokes” of a particular size (width) and orientation (direction). Final denoising is achieved with the construction of gradient maps, thresholding of selected coefficients and reconstruction of an “enhanced/denoised” volume. The reconstructed data serves as an initial guess for volume extraction. The potential of the proposed methodology resides in the mathematical assertion that the problem of 3D surface denoising and detection is more constrained, more accurate in localization (time and space) and more robust when compared to traditional methods of 2D denoising and segmentation.

2 Motivation and Methodology

Raw data acquired from 4D ultrasound has specific characteristics relevant to problems in segmentation. Feature detection should be robust to speckle noise and attenuation artifacts. Speckle noise corrupts the data by introducing sharp changes in an image intensity profile, while attenuation alters the intensity of equally significant cardiac structures, depending on their orientation with respect to the position of the ultrasound beam. *These acquisition properties suggest that measures based on phase information rather than intensity might be more appropriate for denoising and edge detection.* This idea was first expressed and developed on ecocardiography images in a paper by Noble et al. [6, 7]. Their work showed promising results and improvement in the quality of edges detected in the spatio-temporal domain in comparison to the Deriche intensity-based method [8]. Moreover, it is interesting to note that the brushlet decomposition is invariant to intensity and contrast range of an image but depends on the spatial frequency content of a signal. This makes it a very attractive and powerful basis for the analysis of 4D cardiac ultrasound where choosing a single global intensity-based edge threshold is not possible due to position dependent attenuation.

Powerful denoising and segmentation methods have been developed in the context of wavelet analysis [9, 10]. The intuitive principle used in this framework is to consider a domain where features of interest in a signal can be decorrelated from noise allowing for a selective reconstruction of the signal features alone. Wavelet basis functions are characterized by finite time support and localized spectrum in the Fourier domain. Wavelet orthogonal bases are constructed by scaling and translating a single “mother wavelet function”. This framework introduces the notion of shift in time and scale, providing a time-scale representation of a signal via wavelet coefficients. Brushlet functions have similar mathematical properties to wavelet functions. A wavelet scale is analogous to a brushlet “brushstroke” characterized by a specific size and an orientation. A brushlet decomposition is carried out on subquadrants of the Fourier domain, where tiling of the domain is analogous to the shifting of wavelet functions, and strong gradient in the transform domain means a discontinuity of pattern orientation. The coefficients do not rely on pixel intensity but on spatial frequency alone.

An overcomplete scheme is essential for accurate preservation of spatial information (features) in the original data set [11].

Brushlet Basis

Brushlet functions are complex valued and well localized in the frequency domain. Their construction is based on a windowed Fourier transform of the Fourier transform of an image. A windowed Fourier basis of the Fourier plane provides a more flexible angular resolution than wavelet packets [12]. The projection on this orthonormal basis of $L^2(\mathfrak{X})$ provides a decomposition of an image along distinct texture orientations. The initial idea motivating the construction of the brushlet is the desire to build an orthonormal basis of transient functions with good time-frequency localization. For this purpose, windowed complex exponential functions have been used for many years in the context of the sine and cosine transforms. Similarly, Meyer and Coifman introduced a division of the real axis into the subintervals $[a_n, a_{n+1}]$, and a new set of basis functions:

$$u_{j,n}(x) = b_n(x - c_n)e_{j,n}(x) + v(x - a_n)e_{j,n}(2a_n - x) - v(x - a_{n+1})e_{j,n}(2a_{n+1} - x)$$

With c_n equal to the middle value of $[a_n, a_{n+1}]$, b_n and v are two window functions,

and $e_{j,n}$ the complex value exponential defined as: $e_{j,n}(x) = \frac{1}{\sqrt{l_n}} e^{-2i\pi \frac{(x-a_n)}{l_n}}$

The windowing functions b_n and v and the basis function $u_{j,n}$ are displayed in Figure 1.

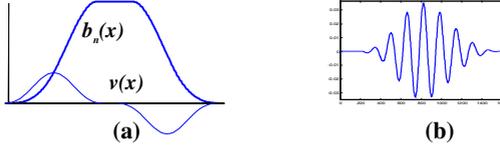


Figure 1: (a), Windowing function b_n , and bump function v defined on the interval $[a_n - \epsilon, a_{n+1} + \epsilon]$. (b), real part of brushlet basis function $u_{j,n}$ defined for $j=5$, $l_n=8$, and $\epsilon=4$.

Suppose we call f a given signal and \hat{f} its Fourier transform. We can project \hat{f} on a brushlet basis, $\hat{f} = \sum_n \sum_j \hat{f}_{n,j} u_{n,j}$ with $u_{n,j}$ the brushlet basis function and $\hat{f}_{n,j}$ the brushlet coefficients. By applying an inverse Fourier transform, we can compute a decomposition of f , $f = \sum_n \sum_j \hat{f}_{n,j} w_{n,j}$ on the orthonormal basis $w_{n,j}$, the inverse Fourier transform of $u_{n,j}$.

The projection of \hat{f} on $u_{n,j}$ is efficiently implemented by a folding technique described by Wickerhauser in [13] and fast Fourier transform (FFT). The reconstruction is simply computed in the same manner by an inverse Fourier transform and unfolding operation.

The 2D and 3D implementations are a direct extrapolation of the 1D projection computation using a tensor product structure. For the 2D brushlet function basis, the x

and y axis are divided into subintervals and on each interval, we define a set of orthonormal 2D basis functions $u_{j,n} \otimes u_{m,k}$ and their inverse Fourier transform pair $w_{j,n} \otimes w_{m,k}$.

An illustration of sample 1D and 2D brushlet basis functions computed in this fashion is shown in Figure 2, below. We can observe the difference in spatial and frequency resolution obtained with two different window sizes of b . A good spatial resolution corresponds to a window with small support and is associated with less frequency resolution as shown in Figure 2(a.2-b.2-c.2). Good frequency resolution corresponds to a window with small support and is thus associated with less spatial resolution as shown in Figure 2(a.1-b.1-c.1).

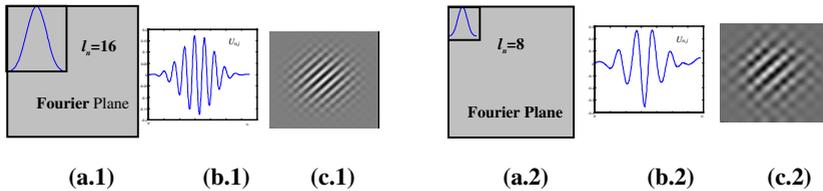


Figure 2: Two-dimensional brushlet basis functions $u_{n,j} \otimes u_{m,k}$ for $j=5, \epsilon=4$, (a) $l_n=16$, (b) $l_n=8$. (a.1-a.2), Selection of size of the quadrants in the Fourier plane for windowing with b . (b.1-b.2), Real part of 1D brushlet basis function. (a.3-b.3), Real part of 2D brushlet basis.

As the number of quadrants increases, the frequency resolution and the number of distinct brushlet orientations represented increases, while the size of the brush stroke and its spatial resolution decrease at the same time. This tradeoff between spatial and frequency resolution is analogous to the Nyquist uncertainty principle for spectral analysis. We illustrate these theoretical properties in Figure 3 with the use of two mathematical phantoms. To best demonstrate the directional selectivity of a brushlet quadrant we chose as a first example, in Figure 3(a.1 -b.1), a non-specific oriented circular pattern with variable spatial frequency in every direction (concentric circular bands of variable size and gray values). The second phantom, in Figure 3(a.2-b.2), is a 2D representation of a chirp signal.

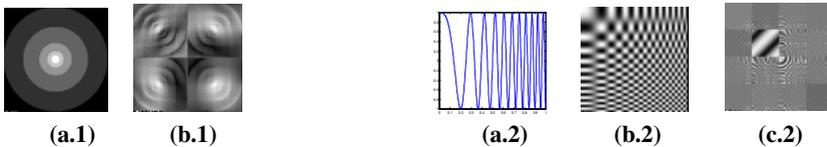


Figure 3: (a.1), A mathematical phantom with concentric bands of variable size and gray value. (a.2), Chirp signal in 1D. (b.2), Cross-product of the 1D signal displayed as an image. (b.1 - c.2), Modulus of the brushlet coefficients for a tiling of the Fourier domain into 4 quadrants (b.1) and 16 quadrants (c.2). Note the different pattern orientations in distinct quadrants.

The implementation of this analysis in three dimensions is straightforward. The tensor product structure of the basis allows a direct extrapolation of the 1D case to both 2D and 3D. In terms of implementation, for a given volume data, its 3D Fourier transform

is first computed, then each sub-block is folded along the three directions and finally expanded into a 3D exponential basis by a 3D Fourier Transform.

Multiscale Analysis

As described above, the brushlet serves as a basis for analysis of the Fourier transform of the original signal and not the signal itself. The number of subintervals along each dimension of the signal (3 for a volumetric data set) and their position in transform space determine the size and the orientation of the brushstroke. The notion of tiling of the Fourier domain of the signal prior to analysis is illustrated in Figure 4.

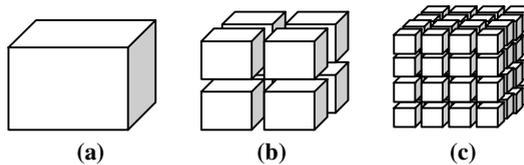


Figure 4: Tiling of the Fourier domain prior to projection on the brushlet basis. (a), Original Fourier Transform of the volume data. (b), Tiling of the Fourier domain into 4 quadrants by dividing each direction in two. Each cube is half the dimension of the original data in each direction. (c), Tiling of the Fourier domain into 64 cubes by dividing each direction by four. Each cube is thus four times smaller than the original volume.

The tiling of the Fourier sub-spaces determines the number distinct directional brushlet basis used for the analysis and their size (i.e., the more quadrants defined, the more *angular directions* for a particular brush stroke are available). The coefficient domain is the same size as the sub-cubes for each brushstroke. Therefore, the smaller the quadrants, the more precise is the frequency resolution of the analysis and the better the resolution at each spatial frequency, or scale. It is important to note that the improvement in frequency resolution implies a loss in spatial resolution as the number of coefficients influencing the reconstruction of the whole signal diminishes with the size of the cube.

The notion of multiscale decomposition can be extended with the brushlet basis by analogy to the wavelet structure. In this context, we analyze our signal with different tiling and manipulate independently different scales for different orientations. Coifman and Meyer exploited this property for optimal compression. In this study, we have used it for denoising and enhancement by eliminating the most outer quadrants as high-frequency noisy components and reconstructing with selected low frequency (inner-quadrants) coefficients. More precisely, in the example displayed in Figure 4, the eight cubes of the tiling in (b) correspond in terms of direction and frequency to the eight inner cubes of the finer tiling shown in (c). The frequency resolution in (b) is less than in (c) but the spatial resolution is inversely finer. Multiscale analysis in this context consists of tracking directional patterns within the eight cubes of (b) corresponding to the eight inner cubes of (c), at a lower scale of analysis.

Overcomplete Framework

Recall that the size of the cubes created with the tiling of the Fourier domain determines the size of the coefficient domain associated with each characteristic brush-stroke of a cube. The diminution of dimension between the spatial domain and the transform domain is analogous to a two-fold downsampling in dyadic wavelet analysis. This downsampling is problematic when manipulating transform coefficients, because there is not a homomorphism between the original signal and coefficient domains.

The theory of overcomplete multiscale analysis has been developed to overcome this mathematical limitation. Overcomplete multiscale representations are well suited for image analysis and denoising/enhancement, because they avoid aliasing effects introduced by critically sampled representations [14] and yield a shift invariant representation. In our case, the aliasing effect arises from the selection of overlapping sub-intervals on the Fourier plane that are expanded into a local Fourier basis. To avoid this and increase the number of coefficients within the same subinterval size, we project onto an extended Fourier basis. This increases the matrix size of the projection domain without changing the original signal. The overcomplete projection is efficiently implemented by padding the original signal with zeros and computing its FFT. Since padding a signal will increase the resolution of the FFT, overcomplete projections increase the number of coefficients for the same interval and therefore the resolution in the transform (coefficient) plane as illustrated in Figure 5.

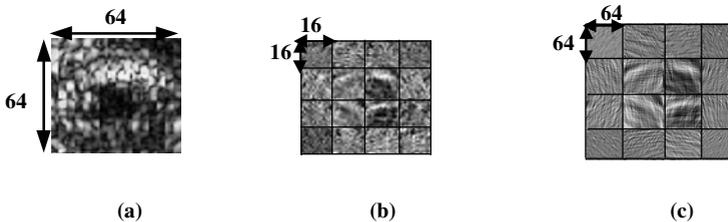


Figure 5: (a) Original cardiac ultrasound slice (64×64 pixels) from 4D data. (b-c), Real part of coefficients for a 16-quadrant tiling of the Fourier plane in the non-overcomplete case (b), and overcomplete case (c). In (b), the dimension of each quadrant is four times smaller than the original image matrix size. In (c), each coefficient quadrant has the same dimension as the original image matrix.

Thresholding for Denoising

Images formed with coherent energy such as ultrasound suffer from speckle noise. This type of noise consists of a granular pattern and is correlated with the surface characteristics of an organ and orientation of the beam. The resulting image degradation has a significant impact on image quality, interpretation and post processing.

More recently several spatial and frequency based denoising techniques have been investigated for echoplanar ultrasound images [15], [16], [17]. However, these traditional methods often reduce noise at the cost of blurring edge features. For this reason multiscale analysis has been intensively tested as a denoising tool [9, 18]. The general

scheme for denoising is achieved by lowering or eliminating noise energy via coefficient thresholding or ‘shrinkage’. Of particular interest to us is the recent work of Zong, Laine and Geiser [19], using non-linear thresholding on dyadic wavelet coefficients for speckle reduction. It showed a significant advantage by reducing inter-physician variability in quantification of cardiac areas, leading to an improved accuracy in the diagnosis of coronary disease. Thresholding is based on the assumption that the noise component can be isolated in certain projection coefficients while energy of a signal is largely concentrated in separate coefficient sets (sub-spaces). This approach has been successfully applied for many years within wavelet decompositions via frames [20], [9].

In our study, we used the threshold selection method for coefficient shrinkage in an orthonormal basis developed by Donoho and Jonhstone [21] and investigated by Mallat in [20]. After projecting original data on a selected basis, we applied a hard thresholding on the projected coefficients in the transform domain as defined below.

The value of the threshold level was empirically deduced from distinct image properties and the projection coefficients’ distribution. After the thresholding operation, the denoised signal was simply reconstructed with the new set of coefficients.

3. Results

Volume analysis with brushlet basis and gradient volume visualization was tested on three different volumes of increasing complexity: (1) a mathematical phantom with three intensity levels of gray, (2) a contrast echocardiogram and, (3) a clinical volume data set.

The 3D-ultrasound acquisition machine used in this study was the Model 1 RT3D (Volumetric Medical Imaging Inc, Durham, NC), originally developed at Duke University [22]. A single volumetric data acquisition cycle requires around 70ms at a depth setting between 10 and 15 cm. A transthoracic 2.5 MHz matrix array transducer scans the 3D volume electronically. The resolution of the acquisition is 64×64 pixels in the short axis plane and an average of 300 pixels in the long axis plane. The clinical volume processed in this study was of size $64 \times 64 \times 64$, and the phantom volumes were made to match these same dimensions.

The three volumes were analyzed with the brushlet basis for a tiling of the Fourier domain in eight cubes with overcomplete representations. This analysis provided a decomposition of four brushstrokes represented by paired diagonal cubes. The brush-stroke orientation in each of the three directions was $\pm 45^\circ$ for each cube.

Visualization in 3D used isosurfaces and isovolumes at selected levels. The level was set as the maximum value of each gradient volume, in the transform domain. The ‘‘marching cube’’ algorithm was used for the isovolume computation [23].

1. A **mathematical phantom** created with three gray levels is illustrated in Figure 6(a.1-b.1-c.1-d.1). It consists of two ovoid shapes, one inside the other one. The initial geometry of the phantom is known *a priori*. The edges were blurred with a median filter, prior to processing. This example is provided to clearly visualize the concept of a gradient volume in the transform domain. The maximum gradient values

were located in the inter-space between the two ovoid objects. Reconstruction after hard thresholding returned the original volume with negligible numerical error.

2. An *in-vitro* **contrast echocardiogram phantom** that consists in a balloon, filled with human albumin and sonicator bubbles as a contrast agent, immersed in a water tank is shown in Figure 6. A volume was recorded with the same 3D-ultrasound machine used in clinical screenings. A typical slice (64×64 pixels) of the volume data is displayed in Figure 6(a.2) below. The inner black cavity corresponds to the albumin tissue inside the balloon. The surrounding white layer is an artifact created by the bubble accumulation on the inner surface of the balloon wall. The corresponding gradient volume in the transform space is displayed in Figure 6(b.2-d.2). We observed a very efficient outline of the inner cavity of albumin from the outside of the balloon.

3. A **clinical volume data** set was acquired with the same clinical 3D-ultrasound transducer. The patient lied on a gantry during the screening examination and an apical view of the left ventricle and atrium is acquired by placing the beam on his/her lateral costal margin. The data are presented in Figure 6(a.3-c.3).

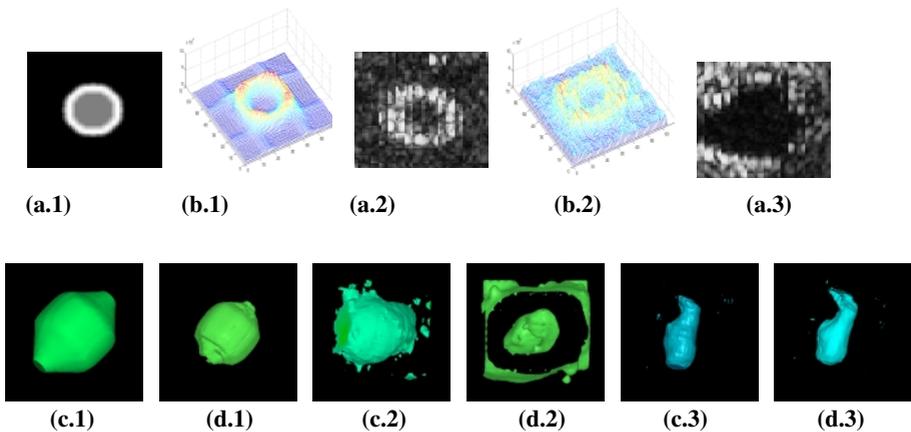


Figure 6: (a.1-b.1-c.1-d.1), Mathematical phantom. (a.2-b.2-c.2-d.2), Contrast echocardiogram phantom. (a.3-c.3-d.3), Clinical volume data. (a.1-a.2-a.3), 2D slice in cross section of the volume data. (c.1-c.2), Volume visualization of the phantoms. (c.1). (b.1-b.2), Plot of the gradient in the transform plane for one slice of the domain. (d.1-d.2), Isosurface of the maximum gradient in the transform domain. (c.3-d.3), Volume visualization of the gradient in the transform plane for an 8-cubes (c.3) and 64-cubes (d.3) tiling of the Fourier domain.

Validation of Gradient volumes in Transform Domain In order to validate our gradient volumes as a “valid” representation of the data in terms of feature shapes and locations we asked an expert in echocardiography to manually digitize the volume data of the contrast phantom and the clinical echocardiograms. The manual segmentation was performed on sixty four 2D slices of the volumes ($64 \times 64 \times 64$ pixels). The result of the superposition of the gradient volumes and the manually segmented volumes is displayed in Figure 7. Regarding the phantom data in Figure 7(a), the match between the gradient data (isosurface of maximum value) in the transform domain and

the manually segmented volume of the albumin cavity is excellent. Please note that the isosurface describes the inner albumin cavity and the outer balloon wall separated by the bubble layer. Indeed the gradient in directional information between the two layers is high. The match for the clinical data, as displayed in Figure 7(b), is also very impressive. The most striking result in the transform domain is our *ability to isolate information from the ventricular papillary muscles and the wall muscular pattern*. The sample manual segmentation did not trace the papillary muscles. This difference in contour is distinguishable in Figure 7(b) at the “neck” of the gradient volume, where the papillary muscles insert (see arrow). This is a remarkable result since systematic errors usually prohibit visualization of papillary muscles in 3D from traditional 2D gated data.

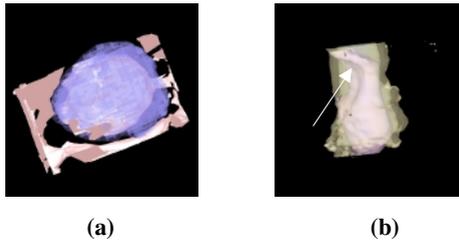


Figure 7: Comparison between manually segmented contours and gradient isosurfaces at maximal values for echocardiographic phantom (a) and clinical data (b). (a), Volume visualization of the gradient inner isosurface surrounded by the manually extracted volume of the balloon. (b), Volume visualization of the gradient isosurface at maximum value surrounded by manually extracted volume of the left ventricle. The papillary muscles were not segmented manually but are correctly isolated by the gradient surface.

Denosing Performance. Denosing performance via hard thresholding was evaluated for an empirical threshold value of 20% of the maximum coefficient values of each cube. The clinical data set was decomposed into an eight-cube ($2 \times 2 \times 2$) tiling of the Fourier domain. Quantitative measurements of image quality improvement are difficult. We plan to follow the validation approach of Laine et al. [19] by using cardiologist performance for denosing quality assessment. A first observation of the denoised volume data showed improvement in the smoothness of the LV edges. However a loss in contrast was observed in the reconstructed voxels.

4. Conclusion

Our method for feature extraction used a hybrid technique to combine model based and directional denosing and segmentation in three dimensions by identifying efficient projection coefficients within sets of redundant articulated (orientation rich) bases. The results of our method will be used to clinically evaluate new 4D ultrasound (3D plus time) acquisition techniques and provide an accurate method to quantify heart volumes from this type of data. This is a very exciting challenge since this type of data is totally new. A future study will aim at modeling the heart volume at systole

and diastole and then study of myocardium deformation patterns to characterize and isolate areas of ischemic tissue. The development of non-invasive measurement techniques of cardiac tissue stress and strain is of major clinical importance. The result of our proposed analysis would provide a valuable imaging tool for modeling cardiac volumes.

References

- [1] J. Duncan, R. Owen, L. Staib, and P. Anandan, "Measurement of non-rigid motion using contour shape descriptors," presented at Computer Vision and Pattern Recognition, 1991.
- [2] Eric Bardinet, Laurent D. Cohen, and N. Ayache, "Tracking and motion analysis of the left ventricle with deformable superquadrics," *Medical Image Analysis*, vol. 1, pp. 129-149, 1996.
- [3] S. Benayoum, C. Nastar, and N. Ayache, "Dense non-rigid motion estimation in sequence of 3D images using differential constraints," presented at Conference on Computer Vision, Virtual Reality and Robotics in Medicine, Nice, France, 1995.
- [4] P. Shi, G. Robinson, C. Constable, A. Sinusas, and J. Duncan, "A model-based integrated approach to track myocardial deformation using displacement and velocity constraints.," presented at IEEE International Conference on Computer Vision, Cambridge, MA, 1995.
- [5] F. Meyer and R. R. Coifman, "Brushlets: A tool for directional image analysis and image compression," *Applied and computational harmonic analysis*, vol. 4, pp. 147-187, 1997.
- [6] Miguel Mulet-Parada and J. A. Noble, "2D+T acoustic boundary detection in ecocardiography," presented at Medical Image Computing and Computer-Assisted Intervention-MICCAI98, Boston, MA, 1998.
- [7] G. Jacob, J. A. Noble, M. Mulet-Parada, and A. Blake, "Evaluating a robust contour tracker on echocardiographic sequences," *Medical Image Analysis*, vol. 3, pp. 63-75, 1999.
- [8] R. Deriche, "Fast algorithms for low-level vision.," *Pattern Analysis and Machine Intelligence*, vol. 12, pp. 78-87, 1990.
- [9] D. L. Donoho and I. Jonhstone, "Ideal denoising in an orthonormal basis chosen from a library of bases," *Comptes Rendus de l'Academie des Sciences, Paris, Serie I*, vol. 319, pp. 1317-1322, 1994.
- [10] J. Fan and A. F. Laine, "Multiscale contrast enhancement and denoising in digital radiographs," in *Wavelets in Medicine and Biology*, A. Aldroubi and M. Unser, Eds. Boca Raton, FL: CRC Press, 1996, pp. 163-189.
- [11] A. F. Laine, J. Fan, and W. Yang, "Wavelets for contrast enhancement of digital mammography," *IEEE Engineering in Medicine and Biology Society Magazine*, vol. 14, pp. 536-550, 1995.
- [12] M. V. Wickerhauser, "Smooth localized orthonormal bases," *Comptes Rendus de l'Academie des Sciences, Paris I*, pp. 423-427, 1993.
- [13] P. Ausher, G. Weiss, and M. V. Wickerhauser, "Local sine and cosine bases of Coifman and Meyer and the construction of smooth wavelets," in *Wavelets- A tutorial in Theory and Applications*, vol. 2, *Wavelet Analysis and its Applications*, C. K. Chui, Ed. San Diego: Academic Press, 1992, pp. 237-256.
- [14] S. Mallat, "A theory for multiresolution signal decomposition: The wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, pp. 674-693, 1989.
- [15] E. A. Geiser, D. C. Wilson, G. L. Gibby, J. Billet, and D. A. Conetta, "A method for evaluation of enhancement operations in two-dimensional echocardiographic images," *Journal of the American Society of Echocardiography*, vol. 4, pp. 235-246, 1991.
- [16] S. M. Collins, D. J. Skorton, E. A. Geiser, J. A. Nichols, D. A. Conetta, N. G. Pandian, and R. E. Kerber, "Computer-assisted edge detection in two-dimensional echocardiography: comparison with anatomic data," *Journal of the American Society of Echocardiography*, vol. 53, pp. 1380-1387, 1984.
- [17] M. J. Lester, J. F. Brenner, and W. D. Selles, "Local transforms for biomedical image analysis," *Computer Graphics, Image Processing*, vol. 13, pp. 17-30, 1980.
- [18] D. L. Donoho and I. M. Johnstone, "Threshold selection for wavelet shrinkage of noisy data," presented at 16th Annual Int. Conference of the IEEE Engineering in Medicine and Biology Society, 1994.

- [19] Xuli Zong, Andrew F. Laine, and E. A. Geiser, "Speckle reduction and contrast enhancement of echocardiograms via multiscale nonlinear processing," *IEEE Transactions on Medical Imaging*, vol. 17, pp. 532-540, 1998.
- [20] S. Mallat, *A Wavelet Tour of Signal Processing*, 1998.
- [21] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation via wavelet shrinkage," *Biometrika*, vol. 81, pp. 425-455, 1994.
- [22] S. W. S. Olaf T. Von Ramm, "Real time volumetric ultrasound imaging system," *Journal of Digital Imaging*, vol. 3, pp. 261-266, 1990.
- [23] A. V. S. Inc., "AVS/Express, Developer Edition," , 4.1 ed. Waltham, MA, 1998.