

Automatic Mitral Valve Inflow Measurements from Doppler Echocardiography

JinHyeong Park¹, S. Kevin Zhou¹, John Jackson², and Dorin Comaniciu¹

¹ Integrated Data Systems, Siemens Corporate Research, Inc., Princeton, NJ, USA
{jin-hyeong.park, shaohua.zhou, dorin.comaniciu}@siemens.com

² Ultrasound Division, Siemens Medical Solution, Mountain View, CA, USA
jacksonjohn@siemens.com

Abstract. Doppler echocardiography is widely used for functional assessment of heart valves such as mitral valve. In current clinical work flow, to extract Doppler measurements, the envelopes of acquired Doppler spectra are manually traced. We propose a robust algorithm for automatically tracing the envelopes of mitral valve inflow Doppler spectra, which exhibit a large amount of variations in envelope shape and image appearance due to various disease conditions, patient/sonographer/instrument differences, etc. The algorithm is learning-based and capable of fully automatic detection and segmentation of the mitral inflow structures. Experiments show that the algorithm, running within one second, yields comparable performance to experts.

1 Introduction

Doppler echocardiography is widely used in clinical practices to assess the heart valve functionality as it records the blood velocity [1][2]. The current work flow of Doppler analysis requires manual tracing the envelopes of acquired Doppler spectra, based on which clinically relevant measurements are computed. The manual tracing is a main bottle-neck of the work flow. In the paper, we aim to propose an automatic algorithm for tracing the envelopes of the Doppler spectra belonging to mitral valve inflow (MI) only.

The MI patterns and measurements have been studied extensively as indices of left ventricular diastolic function [2]. Fig. 1 displays a few sample images of MI spectra with the expert envelopes overlaid to illustrate the challenges we confront. The MI pattern, occurring in the ventricular diastole phase, typically consists of an “early” wave (E-wave) and an “atrial” wave (A-wave). For normal hearts, the E- and A-waves do not overlap each other and the E-wave is higher than the A-wave. But for disease hearts, the following can occur: the E- and A-waves sometimes overlap depending on heart diseases; the E-wave is lower than or of the same height as the A-wave; only the E-wave is present with no A-wave. The above factors contribute significant variation in the envelope shape. In terms of image appearance, the variation is large too due to signal aliasing, difference in imaging setting, etc.




Fig. 1. Examples of MI Doppler spectra. (a) Isolated triangles. (b) Overlapping triangles.

There are a few previous approaches that address the automatic tracing of Doppler spectra envelopes. In [3][4], image processing/filtering techniques such as low-pass filtering, thresholding and edge detection are used. They, however, do not guarantee robustness in the presence of severe images artifacts. Recently a learning-based method, the so-called *PHD* framework [5], was proposed for detection and segmentation of deformable anatomic structures in medical images. This method was successfully applied for automatically tracing the envelopes of three Doppler flow types including MI. In [5], a triangle representation was used to model the MI envelope. This representation has inherent difficulty when dealing with severely overlapping E- and A-waves, rendering the missing triangle roots, which are often seen in the spectra of diseased hearts.

We propose a robust algorithm for automatically tracing the envelopes of mitral valve inflow Doppler spectra. The proposed algorithm builds upon the techniques used in the *PHD* framework [5]. However, unlike [5], it explicitly handles the overlapping E- and A-waves by separate detection and segmentation of non-overlapping E/A-wave and overlapping E/A-waves. The non-overlapping case is formulated as a problem of single triangle detection/segmentation, which means that the E- and A-waves are treated as the same object with a triangle shape. For the overlapping case, we detect and segment a pair of overlapping E- and A-waves. In addition, we introduce several novel components: direct shape inference from image appearance for the overlapping structures, shape score computation based on image gradient, information fusion for robust decision making, etc.

2 Automatic Computation of MI Measurements

To achieve automatic detection and segmentation of target deformable objects, the proposed algorithm first employs a series of detectors to anchor the location of the objects and then utilizes the shape information to segment the deformable object. Fig. 2 gives the schematic overview of the algorithm.

A single triangle object represents an isolated E-wave or A-wave, and a double triangle object represents a pair of overlapping E- and A-waves. A single triangle




Fig. 2. Algorithm overview. (a) input image, (b) results of 1st layer detectors, (c) results of 2nd layer detectors, (d) results of shape inference, (e) Best candidates of single triangle and overlapped triangles, and (f) Final segmentation results.

model consists of three points: a left root, a right root and a peak; a double triangle model has five points: a left root, a right root, a left peak, a right peak, and an intersection point. Base on the observation that both single and double triangle objects share left root, right root, and peak(s), we train three part detectors: left root detector (LRD), right root detector (RRD), and peak detector (PKD). Note that the double triangle object has two peaks but we pool them together during training. We also train two global box detectors: a single triangle detector (STD) and a double triangle detector (DTD). These detectors are organized in a hierarchical manner to efficiently prune the search space down to a set of most promising candidates for both single triangle and double triangle objects. Each candidate is associated with a posterior detection probability $P_D(\Theta|I)$, where Θ is object parameter and the I denotes the image input. Refer to Section 2.1 and [5] for more detail about how they are organized and collaborated.

Once a candidate box for a single triangle is identified, the triangle shape is fully specified. But, this is not the case for a double triangle candidate box as the intersection point of the two overlapping triangles is unknown. We therefore insert a shape inference model for the double triangle object. For each shape candidate, either specified or inferred, the shape profile model is invoked to score it with a posterior probability $P_S(\Phi|I)$, where Φ is the shape model. Section 2.2 addresses how to compute $P_S(\Phi|I)$.

Based on the probabilities $P_D(\Theta|I)$ and $P_S(\Phi|I)$, the algorithm selects the best candidates from the single triangle candidate pool and the best double triangle candidate. Then we compare them to determine the final result.

Fig. 2(a-f) depicts the results of important steps in the algorithm. In Fig. 2(b), yellow dots, red dots, and green dots represent the detection results of LRD, RRD and PKD respectively. In (c), the white boxes are the detection results of STD and the green boxes are those of DTD. The shape inference results computed from the detection boxes are plotted in (d). As shown in (c) and (d), several candidate results around each structure are produced. In (e), it shows the results the algorithm selects the best candidates for STD and DTD independently from their candidate pool. In (f), the final segmentation results, obtained by comparing the better best candidates in (e), are depicted. Once the segmentation is done, measurements can be computed. Sections 2.1 - 2.3 will present the detailed algorithm.

2.1 The Detection Probability

As earlier mentioned, we train one set of LRD, RRD, and PKD for both single triangle and double triangle structures. For box detector, we collect separate training data for single triangle structures and double triangle structures and train the STD and DTD. We treat the detection problem as a two-class classification problem (positive vs. negative) and follow the probabilistic boosting network implementation [6] for testing efficiency. Each detection carries a detection probability. For example, the LRD outputs the posterior probability of being the left root object O_{LR} given an input image I and a hypothesized location θ_{LR} , which is denoted by $P(O_{LR}|I, \theta_{LR})$.

All the above detectors are organized into a two-layer hierarchy. The first layer consists of LRD, RRD, and PKD, which outputs independent candidates of the left root, right root, and peak. Note that in [5], there are no LRD, RRD trained for the MI type. Using LRD and RRD, we significantly speed up the processing time as they are simple and with light computation yet effectively prune the search space.

The second layer consists of STD and DTD. The STD further verifies if a particular combination of a left root, a right root and a peak point from the candidates form a valid E-wave or A-wave. If the combination passes the STD, the two root points and the peak point form a trace of a triangle envelope and no further segmentation process is needed. On the other hand, the DTD verifies if a combination of a left root, a right root and two peak points that lie in between the left and right roots forms a valid pattern. Even if the pattern is valid, however, we only have four points with the intersection point of E-wave deceleration line and the A-wave acceleration line missing. We need to estimate its location, which will be discussed in the next section.

A target object O parameterized by Θ consists of M parts $O = \{O_1, O_2, \dots, O_M\}$ with part O_i parameterized by θ_i . In our case, we define two target objects: (i) a single triangle object with a left root (LR), a right root (RR) and a peak point (PK), i.e., $O_{ST} = \{O_{LR}, O_{RR}, O_{PK}\}$ and $\Theta_{ST} = \{\theta_{LR}, \theta_{RR}, \theta_{PK}\}$, and (ii) a double triangle object with a LR, a RR, an E-wave peak point (EPK), and an A-wave peak point (APK), i.e., $O_{DT} = \{O_{LR}, O_{RR}, O_{EPK}, O_{APK}\}$ and $\Theta_{DT} = \{\theta_{LR}, \theta_{RR}, \theta_{EPK}, \theta_{APK}\}$.

The *PHD* framework assumes conditional independence among the parts and the global structure. For the single triangle object, it evaluates the object detection probability as

$$P_D(\Theta_{ST}|I) \equiv \prod_{a \in A} P(O_a|I, \theta_a) P(O_{ST}|I, \Theta_{ST}), \quad (1)$$

where $A = \{LR, RR, PK\}$. A similar detection probability $P_D(\Theta_{DT}|I)$ can be define for the double triangle object. The conditional independence assumption brings computational advantage: If one of the classifiers fails, the overall detection fails.

2.2 The Shape Probability

For the double triangle structure, we need to estimate the missing intersection point of E-wave deceleration line and the A-wave acceleration line. To this end, we employed the shape inference algorithm proposed in [7].

Given training images and their corresponding shapes, we attempt to learns a nonparametric regression function that gives a mapping from an image to its shape. The training of shape inference model conducts a feature selection process. Each image is represented by an over-complete set of features. The training data are first clustered in the shape space into several clusters. The algorithm then selects a small set of features from the huge feature pool based on a forward feature framework by maximizing the Fisher separation criterion of the clusters. After training, a training image I_j , whose shape is Φ_j , is represented by a feature vector f_j . Given a query image I , we first compute its corresponding feature vector f , then invoke the nonparametric regression function to infer the shape s :

$$\Phi = \frac{\sum_j \Phi_j k(f_j, f)}{\sum_j k(f_j, f)}, \quad (2)$$

where k is a kernel function.

To define the shape probability $P(\Phi|I)$, we use the image evidence along the shape. Suppose that the trace of the envelope has N discrete points, $\{s_i = (x_i, y_i)\}_{i=1}^N$, distributed with equal distance along the trace. For each point s_i , we compute a shape profile score ψ_i based on the intensity gradients along the trace of the envelope. Let line l_i be perpendicular to the tangent of a point s_i . We denote points on the line l_i by $\{s_{ij} = (x_{ij}, y_{ij})\}_{j=i-\Delta}^{i+\Delta}$, where $\Delta > 0$ and $s_i = s_{ii}$, and assume that $\{s_{ij} = (x_{ij}, y_{ij})\}_{j=i-1}^{i-\Delta}$ corresponds to the outside of the contour and $\{s_{ij} = (x_{ij}, y_{ij})\}_{j=i+1}^{i+\Delta}$ the inside. Then, i^{th} shape profile score is defined as follow:

$$\psi_i = \sum_{j=i}^{i+\Delta} I(s_{ij}) - \sum_{j=i-\Delta}^{j=i} I(s_{ij}), \quad (3)$$

where $I(s_{ij})$ represents a pixel intensity at the location of s_{ij} . The shape probability, $P_S(\Phi|I)$ is then defined using a Sigmoid function:

$$P_S(\Phi|I) \equiv [1 + \exp(-\gamma \sum_i \psi_i)]^{-1}, \quad (4)$$

where $\gamma > 0$ is a pre-specified constant.

By integrating both the detection posterior probability $P_D(\Theta|I)$ and shape posterior probability $P_S(\Phi|I)$, we select the best candidates of single triangle and double triangle per each heart cycle among a cluster of detection results. Again we assume that the detectors and shape profiles are independent of each other, which is reasonable because they are two heterogeneous models, we obtain the fused probability as

$$P(\Theta, \Phi|I) = P_D(\Theta|I)P_S(\Phi|I), \quad (5)$$

Finally we select the best candidates that locally maximizes the fused probability.

2.3 Measurement Computation

For the single triangle model, after selecting one best candidates from the cloud of candidates around a structure, we need to determine whether it is E-wave or A-wave in order to compute the necessary measurements. We rely on the given End of Systole (ES) and End of Diastole (ED) lines: the E-wave appears first and the A-wave follows in the diastole period.

For final decision to take the best solution from the single triangle model or that from the double triangle model, we only rely on the shape probability because the detection probabilities from the two models are heterogeneous.

Four MI measurements [1][2] are computed: E-Wave Peak Velocity (EPV), E-Wave deceleration time (EDT), A-Wave Peak Velocity (APV), and A-Wave duration (ADU). Fig. 3 illustrate the measuring process.




Fig. 3. Four measurements: E-Wave Peak Velocity (EPV), E-Wave deceleration time (EDT), A-Wave Peak Velocity (APV), and A-Wave duration (ADU)

3 Experimental Results

We collected 255 Dicom files as training data and 43 Dicom files as test data, which have several cardiac cycles (3.8 cycles/file). The training data were annotated by a sonographer, and the test data were annotated by two sonographers. To best of our knowledge, this is the largest study so far in the literature. In [5], 153 images were used for training and 46 for testing. Also, it faced significant difficulty in dealing with severely overlapping E/A-waves.

For each cardiac cycle, we computed the consensus ground truth (GT) measurements by averaging the measurements individually computed using the annotations by the two experts. The GT measurements are compared with those

Table 1. comparison of correlation coefficient (CC), structure overlap ratio (OR), and delta Peak Velocity (PV) using test data set

	Correlation Coefficient				OR	$ dPV $ (cm/s)
	EPV	EDT	APV	ADU		
Algorithm vs. Expert Ave.	0.987	0.821	0.986	0.481	85.2%	4.27
Expert 1 vs. Expert 2	0.985	0.903	0.973	0.767	85.8%	3.82




Fig. 4. Scatter plots of the four measurements. The 1st row: the algorithm vs. the average of the experts. The 2nd row: expert 1 vs. expert 2. The columns from left to right: EPV, EDT, APV and ADU.




Fig. 5. Segmentation results

computed by our algorithm. We used correlation coefficient (CC) as a comparison metric. The CC between the two experts is also computed. Table 1 shows the comparison of CC between our algorithm vs. GT and that between the two experts, and Fig. 4 displays the scatter plot of the measurement results.

As shown in Table 1 and Fig. 4, our algorithm performs favorably for EPV and APV measurements when compared to the inter-expert variation. The automatic performance of EDT and ADU needs some improvement to match the expert level. To gauge the detection accuracy, we also measure the standard area overlapping

ratio (OR), i.e., $OR = 2 * area(A \cup B) / (area(A) + area(B))$ and the absolute peak velocity difference ($|dPV|$). Again, the automatic algorithm yields results with the inter-user variability. The algorithm, implemented on a 2GHZ PC with 2GB RAM, runs very fast: it takes less than one second to process one image.

Fig. 5 presents the automatic tracing results on various images. In Fig. 5, the red line represents ground truth annotated by an expert and the green line corresponds automatic tracing computed by the proposed algorithm.

4 Conclusion

We have proposed a robust algorithm for automatic Doppler measurement of MI pattern in spite of significant variations in shape and appearance. To handle severely overlapping of E-wave and A-wave, the algorithm has trained two global structure detectors, one for non-overlapping structures and the other for overlapping structures. We have also proposed to use shape profile to compute shape probability, and implemented a novel framework to combine the detection results and the shape score. The experimental results show that the algorithm performs very comparable to experts while running within one second.

References

1. Feigenbaum, H., Armstrong, W., Ryan, T.: Feigenbaum's Echocardiography. Lippincott Williams & Wilkins (2005)
2. Goldberg, B., McGahan, J.: Atlas of Ultrasound Measurements. Mosby Elsevier, Amsterdam (2006)
3. Tschirren, J., Lauer, R., Sonka, M.: Automated analysis of doppler ultrasound velocity flow diagrams. *IEEE Transactions on Medical Imaging* 20, 1422–1425 (2001)
4. Greenspan, H., Shechner, O., Scheinowitz, M., Feinberg, M.: Doppler echocardiography flow-velocity image analysis for patients with atrial fibrillation. *Ultrasound in Medicine and Biology* 31(8), 1031–1040 (2005)
5. Zhou, S.K., Guo, F., Park, J., Carneiro, G., Jackson, J., Brendel, M., Simopoulos, J., Otsuki, J., Comaniciu, D.: A probabilistic, hierarchical, and discriminant (phd) framework for rapid and accurate detection of deformable anatomic structure. In: *Proc. of ICCV* (2007)
6. Zhang, J., Zhou, S., Comaniciu, D.: Joint real-time object detection and pose estimation using probabilistic boosting network. In: *Proc. of CVPR* (2007)
7. Georgescu, B., Zhou, X.S., Comaniciu, D., Gupta, A.: Database-guided segmentation of anatomical structures with complex appearance. In: *Proc. of CVPR*, pp. 429–436 (2005)