

Segmenting the Left Ventricle in 3D Using a Coupled ASM and a Learned Non-Rigid Spatial Model

Stephen O'Brien, Ovidiu Ghita, and Paul F. Whelan

Centre for Image Processing and Analysis,
Dublin City University, Dublin 9, Ireland

stephen.obrien@eeng.dcu.ie

Abstract. This paper presents a new approach to higher dimensional segmentation. We present an extended Active Shape Model (ASM) formulation for the segmentation of multi-contour anatomical structures. We employ coupling and weighting schemes to improve the robustness of ASM segmentation. 3D segmentation is achieved through propagation of a 2D ASM using a learned non-rigid spatial model. This approach does not suffer from the training and aligning difficulties faced by direct 3D model-based methods. Experimental results are encouraging at this early stage, and future directions of research are provided.

Keywords: segmentation, coupled ASM, adaptive weighting, active shape model, left-ventricle, non-rigid model.

1 Introduction

Cardiovascular disease (CVD) accounted for 17.5 million deaths worldwide in 2005. If trends continue, CVD will claim 20 million lives, annually, by 2015 [1]. In 2008 CVD was responsible for approximately half of all deaths in Europe. Today, modern MRI scanners produce 3D+time images of the heart for diagnosis of CVD. Deriving quantitative measurements of the left ventricle from these images, such as ejection fraction, cardiac output and myocardial thickening, provides crucial indicators of cardiac health and therefore plays a central role in CVD diagnosis. However, manual extraction of left ventricle contours is an especially difficult task due to intensity variations, poor contrast between the epicardium and surrounding thoracic cavity and the presence of papillary muscles which can cause diffuse endocardium borders in 2D images. Automated image segmentation techniques are a promising alternative to traditional manual methods that avoid: operator fatigue, inter/intra observer variability and providing repeatable results, crucial in long-term patient monitoring.

In recent years, it has become widely recognised that automatic segmentation methods need to maximise *a priori* knowledge of the segmentation target to

be effective[2]. Model-based methods have been shown to perform well in noisy medical images, where anatomical structures are often poorly defined. The Active Shape Model (ASM), introduced by Cootes et al.[3], is one such popular model-based approach, where shape variation is captured in linear models using Principal Component Analysis (PCA). ASMs, and their later extension to include texture variation Active Appearance Models (AAM), have been applied to a range of medical applications, including the brain[4], lung [5] and cardiac[6, 7] segmentation. Early ASM and AAM implementations were primarily developed to address the segmentation of 2D data. In recent years there has been a growing trend towards 3D and 3D+time medical image segmentation[8–10].

However, generalising the ASM to 3D (and 3D+time) faces many problems. Specifically, to capture the full variation of a 3D shape requires significantly more training samples and landmarks than the 2D case. Aligning the training set poses a difficult challenge since one has to deal with more degrees of freedom. Finally, 3D ASM initialisation and optimisation are computationally intensive and often become trapped in local minima. Many publications have dealt with the issues with direct 3D ASM implementations, especially constructing 3D ASMs in the presence of small training sets [4]. We propose a new approach to higher dimensional modelling and segmentation that avoids the problems outlined above. The main contributions of this work lies in combining a learned non-rigid spatial model and an extended 2D ASM formulation to achieve 3D segmentation.

2 Method

2.1 ASM Theory

When building an ASM from a set of training shapes, the segmentation target must be descretised into N landmarks. A shape vector, \mathbf{x} , containing N landmarks is described as:

$$\mathbf{x} = [x_0, x_1, \dots, x_N] \quad (1)$$

where x_i is the i^{th} landmark. The training set containing M training shapes is defined as:

$$T = [\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_M] \quad (2)$$

After the training set (T) has been aligned (T'), PCA is applied to extract the eigenvectors which describe the shape variation of T' . To align the shapes in the training set, only translation effects need to be removed as we wish to capture the scale variation of the training samples. Rotational effects are negligible since the contours are all taken from a specific short axis slice and temporal volume. Therefore, the training set can be aligned without applying Procrustes analysis [3]. The most significant eigenvectors Φ are chosen to capture 95% of the variation in the training set [3]. New shapes can be synthesised using Φ and the mean-shape of T' , $\bar{\mathbf{x}}$, according to:

$$\mathbf{x} \approx \bar{\mathbf{x}} + \Phi \beta \quad (3)$$

where β represents the model parameters and \mathbf{x} is the shape vector (1). The limits on the i^{th} component of the β vector is constrained by $\pm 3\sqrt{\Phi_i}$, where Φ_i describes the variance of the i^{th} element of β in T' [3].

Fitting the trained ASM to an unseen image is achieved with an iterative two-stage process. First, the energy of each landmark is minimised according to the local image information (usually image gradient or local texture). This is a data-driven process where the search space for each landmark is constrained along a 1D profile, normal to the direction of the contour. For example, the i^{th} landmark x_i has a corresponding grey-scale profile $L(x_i)$:

$$L(x_i) = \{I(x_i - j\vec{n})\}, -k \leq j \leq k \quad (4)$$

where $2k$ is the length of the profile, \vec{n} is the normal direction of the contour at the landmark x_i and $I(x_i)$ is the grey-scale value in the image at the landmark x_i . In the standard ASM formulation, the minimised landmark x'_i is chosen according to the maximum of the gradient profile:

$$x'_i = \max(\nabla L(x_i)) \quad (5)$$

where $\nabla L(x_i)$ is the gradient profile of $L(x_i)$. Then, the ASM approximates the set of minimised landmarks, \mathbf{x}' , by synthesising a new shape, \mathbf{x}'' , which is the best fit between the captured variation of the ASM and \mathbf{x}' . This is described by:

$$\mathbf{x}'' \approx \bar{\mathbf{x}} + \Phi \beta' \quad (6)$$

$$\text{where } \beta' = \Phi^T (\mathbf{x}' - \bar{\mathbf{x}}) \quad (7)$$

This process continues until convergence is reached.

2.2 Single Shape Multi-contour ASM

When building an ASM of the left ventricle, the variations of the endo- and epicardium need to be captured. These two contours are strongly related anatomically, and so are treated as a single shape (as opposed to creating a separate ASM for each contour). This is achieved by extending the shape vector (1) to describe both contours as one shape. All subsequent calculations to derive Φ , $\bar{\mathbf{x}}$ etc. are similar to the single contour ASM case, described in Section 2.1.

During landmark optimisation, energy is minimised on an individual basis - i.e. no contextual knowledge is used from surrounding landmarks. This is not a limitation when dealing with a single contour ASM where searching is constrained along contour normals so there can be no overlapping between landmark profiles.

However, when two contours are treated as a single shape landmark profiles *can* overlap. This compromises the stability of the ASM in the case of left ventricle segmentation, due to the proximity of the endo- and epicardium as illustrated in Figure 2.2(a). The endocardium is generally the stronger feature of the two due to the contrast between the bright blood pool and the darker myocardium. Therefore, landmarks for both contours tend to converge on the

endocardium in the image (Figure 2.2(b)). The ASM will be unable to approximate the set of landmarks \mathbf{x}' , since such a case will never arise in the training set. The resulting ASM approximation, \mathbf{x}'' , will have little correspondence to the minimised landmarks \mathbf{x}' and, in fact, can go someway to hiding the underlying problem (Figure 2.2(c)).

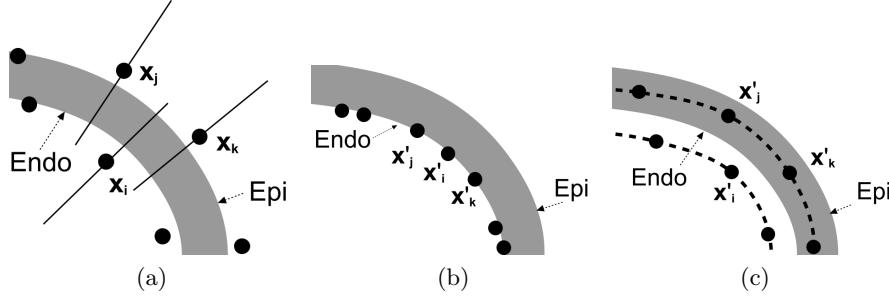


Fig. 1: Sub-figure (a) shows how landmark profiles overlap when minimising the energy of a single shape multi-contour ASM. (b) shows the resulting energy minimisation (assuming the endocardium is the stronger feature in the image). (c) shows the model approximation.

2.3 Proposed ASM Contour Coupling Approach

To address this weakness in the standard ASM (when applied to multi-contour structures) we propose to extend the standard ASM formulation to enforce *a priori* structure during segmentation. Specifically, we enforce biological integrity that forces the epicardium to be outside the endocardium. In the standard approach, each landmark (x_i) is minimised according to the maximum of it's gradient profile $\nabla L(x_i)$ (5). In order to couple the contours, i.e. to enforce *a priori* constraints, we extend (5) as follows:

$$x_i'^{endo} = \max(\nabla L(x_i^{endo})) \quad (8)$$

$$x_i'^{epi} = \max(\nabla L(x_i^{epi}) \setminus \Omega_{endo'}) \quad (9)$$

where $A \setminus B = \{x \in A \mid x \notin B\}$ and $\Omega_{endo'}$ is the region defined by the minimised endocardium landmarks $\mathbf{x}'_{endo'}$. Here, we minimise the energy of the endo- and epi-cardium landmarks differently. We give preference to the endocardium landmarks, which are minimised first, as the endocardium contour is the stronger feature in the image. The profile for each epicardium landmark is excluded from $\Omega_{endo'}$, ensuring structural consistency is maintained during segmentation. This is the central concept of the proposed contour coupling approach and is illustrated in Figure 2.

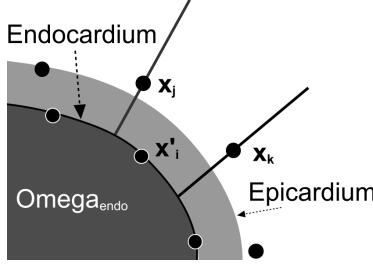


Fig. 2: Contour Coupling. Landmark profiles for x_j and x_k are truncated by $\Omega_{endo'}$ as described in Equation (9)

2.4 Contour Weighting

Sections of the epicardium contour are often poorly defined due to low contrast with surrounding tissues in the thoracic cavity. To improve epicardium segmentation robustness we introduce an adaptive weighting term which reinforces weak edge information. We exploit the knowledge that the endo- and epi-cardium are approximately concentric. Before landmark minimisation, an edge frequency histogram is generated from each epicardium gradient profile, $\nabla L(x_i^{epi})$. The histogram is defined as:

$$H = \frac{1}{N} \sum_{j=1}^N \nabla L(x_j^{epi}) \quad (10)$$

where the epicardium is described by N landmarks. This histogram is used as a weighting term during epicardium landmark minimisation which extends Equation (8) and (9) to:

$$x_i^{endo} = \max(\nabla L(x_i^{endo})) \quad (11)$$

$$x_i'^{epi} = \max((\nabla L(x_i^{epi}) \setminus \Omega_{endo'})H) \quad (12)$$

Here, the epicardium gradient profiles are weighted using the dynamic histogram. The histogram gives weighting to *common* edges outside $\Omega_{endo'}$ region, thus bringing global context into the local landmark minimisation process.

2.5 Propagation to 3D

The common approach to modelling a 3D structure involves the construction of direct 3D shape models [6, 11–13]. Despite the popularity of this approach - it suffers from several flaws. Most importantly, direct modelling of 3D structures is not scalable. This is primarily due to the large increase of degrees of freedom when moving from the 2D to the 3D domain, and further from the 3D to the 3D+time domain. Moreover, large training databases are needed to train 3D, and especially 3D+time models. Finally, the difficulty with aligning the data

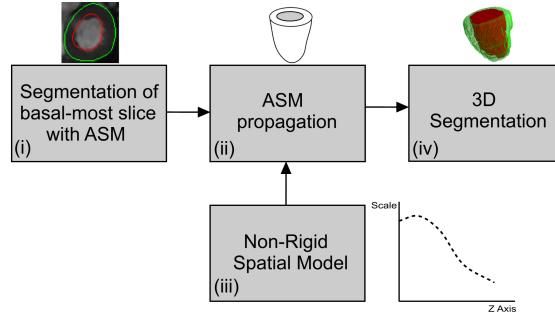


Fig. 3: Flowchart for 3D segmentation algorithm.

during training, as well as the increase of local minima also add to the problems associated with direct 3D and 3D+time modelling.

In this paper, a new scalable method of higher dimensional modelling and segmentation is detailed, when applied to the segmentation of the endo- and epi-cardium in 3D data. This approach does not rely on large databases and does not suffer from aligning or local minima problems. A model is trained that captures the scale variation of the left ventricle over the z axis (Figure 3 (iii)). Scale variations are derived from the manual annotations of the systolic and diastolic cardiac phases from the training set. Scale samples are normalised and used to create a probability density function describing mean scale variations on the z axis.

The initial 2D ASM segmentation of the basal-most slice of the left ventricle (Figure 3 (i)) interacts with the statistically estimated scale model to produce an estimation of the scale of the shape in the next slice. The ASM is initialised with this estimated scale and constrained during segmentation according to the learned variance for that slice from the scale model. Segmentation is propagated in this manner (Figure 3 (ii)) to achieve 3D segmentation (Figure 3 (iv)).

3 Results

The proposed 3D segmentation technique requires two seed points. The first gives both the initial spatial slice of the left ventricle in the dataset and the initialisation for the ASM. The second point defines the final spatial slice of the left ventricle in the volume at the apex. The method is evaluated against 30 3D cine-MRI cardiac images. Both systolic and diastolic cardiac phases are included in the testing database. The 3D scale model is trained from the manual annotations of 30 different 3D cine-MRI cardiac images, containing both normal and pathological patients. Full details of the data and evaluation procedures can be found in Radau et al. [14]. Average segmentation errors are measured at 3.73 mm and 3.16 mm for the endocardium and epicardium, respectively. The Dice measure was also used to detect percentage error in overlap between the segmented region and the manually annotated region. Dice error is measured at 80.9% and

91.3% for the endocardium and epicardium, respectively. The results show that the accuracy of the epicardium is higher than that of the endocardium. This is due, in part, to the effect of the papillary muscles obscuring the true endocardium border during endocardium landmark minimisation. Figure 4 shows a sample volume segmentation using the proposed method.

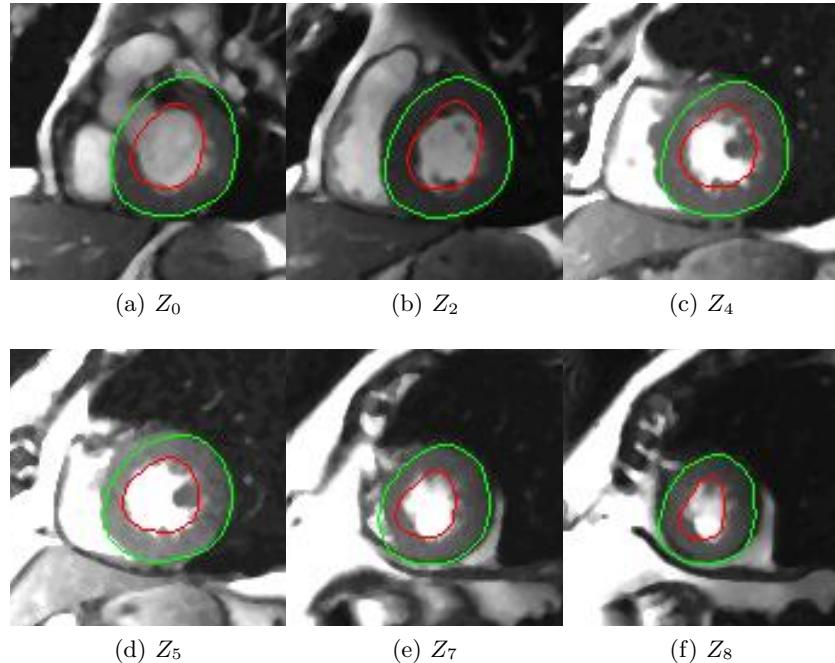


Fig. 4: Sample volume segmentation result using the proposed method. Note: only a partial volume is shown due to space restrictions.

4 Conclusion

This paper describes the implementation and evaluation of a new 3D segmentation technique using a coupled and weighted ASM formulation in conjunction with a learned non-rigid spatial model. The proposed approach addresses the segmentation of the endo- and epi-cardial interfaces in 3D+time cardiac MRI data. This approach is evaluated on real 3D+time cardiac MRI datasets. Experimental results are encouraging at this stage of research. Future directions of research include: improved segmentation accuracy of the 2D ASM, especially in the presence of papillary muscles, and extension of this approach to 3D+time segmentation.

Acknowledgements This work is supported by the Irish Research Council for Science, Engineering and Technology (IRCSET), with additional support provided by the National Biophotonics and Imaging Platform Ireland (NBIP).

References

1. (WHO), W.H.O.: *Cardiovascular disease statistics*. (2004)
2. Frangi, A., Niessen, W., Viergever, M.: Three-dimensional modeling for functional analysis of cardiac images, a review. *IEEE Transactions on Medical Imaging* **20**(1) (2001) 2–5
3. Cootes, T., Taylor, C., Cooper, D., Graham, J., et al.: Active Shape Models—Their Training and Application. *Computer Vision and Image Understanding* **61**(1) (1995) 38–59
4. Zhao, Z., Aylward, S., Teoh, E.: A Novel 3D Partitioned Active Shape Model for Segmentation of Brain MR Images. *Lecture Notes in Computer Science* **3749** (2005) 221
5. Shi, Y., Shen, D.: Hierarchical Shape Statistical Model for Segmentation of Lung Fields in Chest Radiographs. Springer (2008)
6. Andreopoulos, A., Tsotsos, J.: Efficient and generalizable statistical models of shape and appearance for analysis of cardiac MRI. *Medical Image Analysis* (2008)
7. Casero, R., Noble, J.: A novel explicit 2D+ t cyclic shape model applied to echocardiography. In: *Medical image computing and computer-assisted intervention: MICCAI International Conference on Medical Image Computing and Computer-Assisted Intervention*. Volume 11., Springer (2008) 527
8. Lynch, M., Ghita, O., Whelan, P.: Segmentation of the Left Ventricle of the Heart in 3-D+ t MRI Data Using an Optimized Nonrigid Temporal Model. *IEEE Transactions on Medical Imaging* **27**(2) (2008) 195–203
9. Montagnat, J., Delingette, H.: 4D deformable models with temporal constraints: application to 4D cardiac image segmentation. *Medical Image Analysis* **9**(1) (2005) 87–100
10. Stegmann, M., Pedersen, D.: Bi-temporal 3D active appearance models with applications to unsupervised ejection fraction estimation. In: *Proc. SPIE*. Volume 5747.
11. Frangi, A., Rueckert, D., Schnabel, J., Niessen, W.: Automatic construction of multiple-object three-dimensional statistical shape models: application to cardiac modeling. *IEEE Transactions on Medical Imaging* **21**(9) (2002) 1151–1166
12. Fritz, D., Rinck, D., Dillmann, R., Scheuering, M.: Segmentation of the left and right cardiac ventricle using a combined bi-temporal statistical model. *Medical Imaging 2006: Visualization, Image-Guided Procedures, and Display*. Edited by Cleary, Kevin R.; Galloway, Robert L., Jr. *Proceedings of the SPIE*, **6141** (2006) 605–614
13. Zambal, S., Hladuvka, J., Buhler, K.: Improving Segmentation of the Left Ventricle Using a Two-Component Statistical Model. *Lecture Notes In Computer Science* **4190** (2006) 151
14. Radau, P., Lu, Y., Connelly, K., Paul, G., Dick, A., Wright, G.: Evaluation framework for algorithms segmenting short axis cardiac MRI. <http://hdl.handle.net/10380/3070>. The MIDAS Journal - Cardiac MR Left Ventricle Segmentation Challenge