Segmentation of the Right Ventricle Using Diffusion Maps and Markov Random Fields

Oliver Moolan-Feroze[1], Majid Mirmehdi[1],
Mark Hamilton[2], and Chiara Bucciarelli-Ducci[2]

[1] Visual Information Laboratory, University of Bristol, BS8 1UB
[2] Bristol Heart Institute, Bristol, BS2 8HW
oliver.moolan-feroze@bristol.ac.uk, majid@cs.bris.ac.uk

Abstract. Accurate automated segmentation of the right ventricle is difficult due in part to the large shape variation found between patients. We explore the ability of manifold learning based shape models to represent the complexity of shape variation found within an RV dataset as compared to a typical PCA based model. This is empirically evaluated with the manifold model displaying a greater ability to represent complex shapes. Furthermore, we present a combined manifold shape model and Markov Random Field Segmentation framework. The novelty of this method is the iterative generation of targeted shape priors from the manifold using image information and a current estimate of the segmentation; a process that can be seen as a traversal across the manifold. We apply our method to the independently evaluated MICCAI 2012 RV Segmentation Challenge data set. Our method performs similarly or better than the state-of-the-art methods.

1 Introduction

The role of magnetic resonance imaging (MRI) of the left ventricle (LV) in the diagnosis of cardiovascular disease is well established. This has resulted in a large body of research into automated LV segmentation [1], a necessary precursor to the extraction of cardiac parameters. Although research in right ventricular (RV) segmentation is comparatively sparse, recent work [2] has highlighted the importance of RV function to cardiac health. This has resulted in a push within the community to establish accurate and effective methods for RV segmentation and a reflection of this is the 2012 MICCAI challenge on RV segmentation [3].

Accurate RV segmentation is difficult. Even in healthy patients, challenges include inter-patient morphological differences, lack of contrast at cavity borders, obliquity of the tricuspid valve annulus, ventricular trabeculation, and shape variation between apical and basal slices. Responding to these difficulties, researchers have tended towards the inclusion of prior knowledge of the RV to inform segmentation. Often this has taken the form of cardiac atlases or statistical shape models.

Segmentation with cardiac atlases requires the non-rigid registration of a single or multiple expert labellings to the target image; an increased number of

P. Golland et al. (Eds.): MICCAI 2014, Part I, LNCS 8673, pp. 682–689, 2014.

which provide robustness to shape variation. An optimal segmentation is then decided through a system of consensus voting or statistical measures [4]. Zaluga et al. [5] implement a multi-atlas method with a course-to-fine strategy. The prevention of mis-registration is also tackled in [6], where intermediate label results are incorporated into atlas-target registrations to improve alignment. Although proven effective, these methods are slow as a result of the multiple non-rigid registrations required to align the atlases.

The statistical model approach encodes expected shape variation typically using some form of principal component analysis (PCA). Both Mitchell et al. [7], and Ordas et al. [8] construct Point Distribution Models (PDM) of both ventricles before applying them in an active appearance/shape model framework [9]. Forgoing the requirement for corresponding landmarks – a lengthy and error prone process needed for PDMs – Grosgeorge et al. [10] construct their model using PCA of signed distance functions. By generating highly deformed shapes, a static prior map is created that encompasses the extent of the variation found within the training data. Although not applied to RV segmentation, distance functions are used in [11]. Through Expectation-Maximisation, shape priors that are increasingly better fitted to the target image are generated.

Although PCA models have proven effective in many applications, recent research has investigated the ability of shape-based Diffusion Maps [12] – a type of manifold learning method – to represent the intrinsic non-linearity found in many medical datasets [13]. Detailing a method for embedding new shapes into a manifold, and providing a solution to the Diffusion Map pre-image problem, this work has paved a way for the practical use of these models for segmentation. This is evidenced in [14] where a Diffusion Map shape prior is combined with a level-set segmentation with good results.

In this work, we explore the role of a Diffusion Map shape model in application to RV segmentation. We propose that such a model will be better able to represent the complex variations displayed by the RV. The main contribution is in the novel combination of a Diffusion Map shape model with a Markov Random Field (MRF) 2D segmentation framework. Through an iterative method, image data and segmentation results from the MRF are used to generate shapes from the model that increasingly resembles the target image – a process that can be seen as a traversal over a shape manifold. The generated shapes in turn influence the MRF in the form of prior probability maps. In Sect. 2, we detail our method including both MRF formulation and Diffusion Map shape generation. In Sect. 3, we first evaluate the ability of the Diffusion Map model to represent the complex shapes within an RV dataset in comparison to a typical PCA model. This is followed by an assessment of our overall segmentation method when applied to the independently evaluated MICCAI RV Segmentation Challenge dataset.

2 Proposed Method

The proposed segmentation method consists of two steps, iteratively applied until convergence. The first is *Segmentation* (Sect. 2.1) in which we generate an

estimation of the segmentation using an MRF framework combined with a prior probability map generated from a manifold shape model. For the initial iteration, the prior is generated from a mean shape and is aligned to the target using manually placed landmarks. The second step is *Prior Update* (Sect. 2.2) where we combine the current estimate of the segmentation with image information to update the prior through a process of manifold traversal. The method terminates when the difference in segmentation estimation between two iterations falls below a threshold.

Integral to the overall process is the manifold shape model. We define a shape as a signed distance function s. The Diffusion Map is constructed over the set of training shapes $\Gamma = \{s_i\}$ where $i \in 1, \ldots, p$ following the method outlined in [13]. To measure shape similarity, we use the distance proposed in [15],

$$d^2\left(s_i, s_j\right) = \sum_{x \in \Omega_s} \left(H(s_i(x)) - H(s_j(x))\right)^2 \ , \tag{1}$$

where $H(.)$ is the Heaviside function. This benefits from being fast, as well as being positive, symmetric, and obeying the triangle inequality. The result of this process is our shape model - an embedding $\Phi(s) = y,\ y \in \mathrm{R}^m$ where m is the dimensionality of the reduced space.

2.1 Segmentation Using MRFs and Manifold Shape Priors

We model the segmentation as an MRF. The field is a graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ where $\mathcal{V}$ is the set of n image pixels $I = (x_1, \ldots, x_n)$ and $\mathcal{E}$ the edges that connect them with their neighbours. A set of labels $\alpha \in \{\mathrm{O}, \mathrm{B}\}$ represent the object and background classes. The labelling $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ applied over the graph constitutes a segmentation. Each label assignment incurs a cost specified by the energy function

$$E_\omega(I, \boldsymbol{\alpha}) = E_\omega^{\mathrm{d}}(I, \boldsymbol{\alpha}) + E_\omega^{\mathrm{b}}(I, \boldsymbol{\alpha}) + E_\omega^{\mathrm{p}}(I, \boldsymbol{\alpha}, M_\omega) \ . \tag{2}$$

The optimal labelling is when E_ω is at its minimum. The first term E_ω^{d} is a data term which measures the sum of the individual labelling costs at each node in $\mathcal{V}$. This is computed using the local image information and a foreground and background Gaussian Mixture Model (GMM), with 1 and 2 modes respectively

$$E_\omega^{\mathrm{d}}(I, \boldsymbol{\alpha}) = \sum_{i \in \mathcal{V}} -\log p(x_i | \alpha_i, \boldsymbol{\mu}, \boldsymbol{\sigma}) \ , \tag{3}$$

where $P(x_p | \alpha_i, \boldsymbol{\mu}, \boldsymbol{\sigma})$ is the probability of the pixel belonging to either the foreground or background given the GMM parameters $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$.

The second term of the energy function, E_ω^{b}, encourages a smooth labelling over the graph by penalising pairs of nodes (i, j) that have differing label assignments, such that

$$E_\omega^{\mathrm{b}}(I, \boldsymbol{\alpha}) = \sum_{\alpha_i \neq \alpha_j, i,j \in \mathcal{E}} \exp\left(-\frac{(x_i - x_j)^2}{2\sigma_{\mathrm{b}}^2}\right) . \frac{1}{dist(i,j)} \ . \tag{4}$$

σ_{b} is varied to compensate for an expected amount of signal noise in the image and $dist(i,j)$ represents the Euclidean distance between the two nodes.

The final term E^{p}_{ω} incorporates prior shape information into the MRF. A shape generated from the Diffusion Map $\hat{s}$ is transformed into a probabilistic atlas image M_{ω}

$$M_{\omega}(i) = \begin{cases} 1 & \text{if} \quad \hat{s}(i) \leq 0 \\ \exp\left(-\hat{s}(i)/\gamma\right) & \text{if} \quad \hat{s}(i) > 0 \end{cases} \tag{5}$$

where the value γ controls the 'spread' of the influence of the prior outside the boundaries of the zero-level set of $\hat{s}$. The prior term is then defined as

$$E^{\mathrm{p}}_{\omega}(I, \boldsymbol{\alpha}, M_{\omega}) = \sum_{i \in \mathcal{V}} - \begin{cases} \log M_{\omega}(T(i)) & \text{if} \quad \boldsymbol{\alpha}_i = \mathrm{O} \\ \log(1 - M_{\omega}(T(i))) & \text{if} \quad \boldsymbol{\alpha}_i = \mathrm{B} \end{cases} \tag{6}$$

where T is a rigid transform that aligns the model with the target image. The MRF is optimised using the graph cut method of Boykov et al. [16].

2.2 Prior Update through Manifold Traversal

Using the learned manifold and the image data from the segmentation target we present a method of generating patient specific shape priors. This technique of iterative prior generation can be seen as a traversal over the manifold where each iteration produces a shape that increasingly resembles the target. An example of the manifold traversal can be seen in Fig. 1. To drive the traversal, we take the signed distance function of our current estimate of the segmentation s^* and transform it into the space of the model using T^{-1}. This is used to query into the manifold. As Diffusion Maps do not provide a simple way of embedding new data into a learned manifold, we use the Nyström extension [13]. This provides an operator $\hat{\Phi}(T^{-1}(s^*)) = y^*$ where y^* represents the coordinates in the space of the manifold. We take y^* as an estimate of the coordinates that would be produced by embedding the true segmentation and find the nearest neighbours $\mathcal{N}$. The shapes that constitute these neighbours are the templates from which a new shape $\hat{s}$ can be generated. This is done by taking their linear combinations $\hat{s} = \sum_{i \in \mathcal{N}} \theta_i s_i$, where $\theta_i \geq 0$ and $\sum_{i \in \mathcal{N}} \theta_i = 1$. By varying the values in θ we influence the generated shape $\hat{s}$.

To ensure that the generated shape $\hat{s}$ resembles the segmentation target we aim to satisfy two conditions: (a) that it resembles our current 'best guess' at the segmentation s^* and, (b) that it fits to the target image, compensating for inaccuracies in s^*. We encode these two conditions in the energy function

$$E_{\eta}(\hat{s}, s^*, I, T, \lambda_{\eta}) = E^{\mathrm{d}}_{\eta}(\hat{s}, T^{-1}(s^*)) + \lambda_{\eta} E^{\mathrm{h}}_{\eta}(I) \ . \tag{7}$$

The first term is d^2 from (1) and will penalise $\hat{s}$ from deviating from s^*. Similar to [11], the second term is a measure of the entropy inside and outside of the image region enclosed by $T(\hat{s})$, such that

$$E^{\mathrm{h}}_{\eta}(I) = \sum_{x_i \in I} - \left((p_{\mathrm{O}}(x_i) \log p_{\mathrm{O}}(x_i)) + (p_{\mathrm{B}}(x_i) \log p_{\mathrm{B}}(x_i)) \right) \ , \tag{8}$$

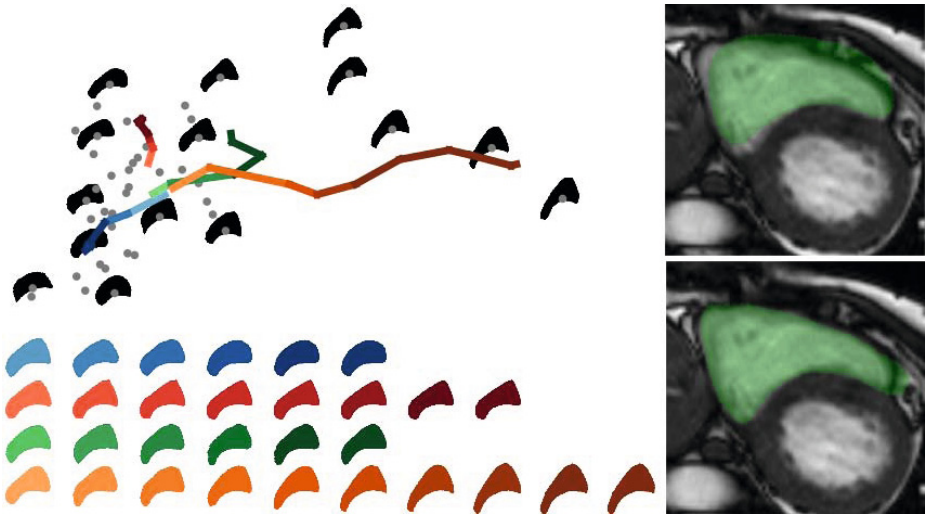

Fig. 1. Visualisation of Manifold Traversal. (Left) Image displaying a learned manifold with traversals overlaid. Black shapes are taken from the training data. Coloured paths map the traversal across the manifold during segmentation. Images below the manifold show the changing shape of the generated prior at each iteration of the segmentation with colours corresponding to their paths. (Right) Example of initial prior (top) and prior at the end of manifold traversal (bottom)

where p_{O} and p_{B} are the probability of image value x_i being found in the foreground or background respectively. These are computed by generating histograms of the image intensity values in the corresponding regions of $T(\hat{s})$. The value λ_η is a weighting term which varies the influence of the entropy calculation. To find the optimal prior, we define the functional $S_{\mathcal{N},\theta} = \sum_{s_i \in \mathcal{N}} \theta_i s_i$ and compute the optimal transform $\hat{T}$ and shape parameters $\hat{\theta}$ using Powell's method [17],

$$\underset{\theta,T}{\operatorname{argmin}} \; E_p(S_{\mathcal{N},\theta}, s^*, I, T, \lambda_\eta) \;\; . \tag{9}$$

3 Evaluation

To evaluate the performance of the proposed method we applied it to the MICCAI 2012 RV Segmentation Challenge dataset [3]. This data was acquired from 32 patients with diagnosed cardiac pathologies, where for each examination, two volumes representing end-diastole (ED) and end-systole (ES) were manually labelled by a cardiac radiologist. The data was split into two sets of 16 examinations, one for training and one for testing. Due to the shape variation between both apical and basal slices and between ED and ES, we split the training data into 11 sets - 6 models at ED and 5 at ES. For all our experiments, the dimensionality of the manifold model was 3.

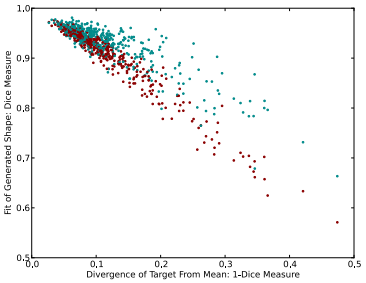

Fig. 2. Plot of manifold (blue) and PCA with 30 modes (red) performance against the divergence of the target from the mean

Table 1. Table showing the DM mean (μ) and standard deviation (σ) from the first experiment

		DM ($\mu \pm \sigma$)
Propsed Method	ED	0.94 ± 0.03
	ES	0.92 ± 0.05
	Total	**0.93 ± 0.04**
PCA *10 modes*	ED	0.93 ± 0.04
	ES	0.88 ± 0.08
	Total	**0.91 ± 0.06**
PCA *30 modes*	ED	0.93 ± 0.04
	ES	0.89 ± 0.08
	Total	**0.91 ± 0.06**

Model Shape Generation Comparison – The first experiment compared the ability of our model against a typical PCA model to represent the complex shapes exhibited by the RV. For both methods, we generated a series of shape models using leave-one-out cross validation across all 11 groupings of training data. The manifold model was built as in Sect. 2, and the PCA models were built by applying PCA to the set of aligned training shapes. The top 10 and 30 modes we retained, which accounted for 97.7% and 99.9% of the variation found within the data. To test our model, we generated a new shape $\hat{s}$ from the test shape s_t by optimising (9) with $\lambda_\eta = 0$ and fixing T as the identity matrix. For the PCA method, we extracted the shape parameters by transforming s_t into the PCA space, regularised them to be within $\pm 3\sigma$ of the learned variation and reconstructed the shape $\hat{s}$ in the typical manner.

Table 1 displays the Dice Metric (DM) over the target s_t and generated shape $\hat{s}$. It shows that all models are able to represent new shapes well, with the manifold model performing slightly better. Also of interest are the results in Fig. 2, where the performance of both our model and the 30 mode PCA model is plotted against the dissimilarity of s_t from the mean shape. This shows that our method is better able to adapt to the extremities of the data variation.

Evaluation on MICCAI 2012 RV Segmentation Dataset – We applied the proposed method to segmentation using the MICCAI RV challenge dataset. We learned the manifold models as described previously. In addition, GMMs were generated for each model to capture the intensity variations between the basal and apical slices. Together, these were applied to the 16 test-set examinations. The results generated by our method were independently evaluated by the organisers of the challenge. As in [10], minimal manual input provided an initial alignment for the model. This consisted of two landmarks at the junction of the RV and LV. Since we optimise the transform T during segmentation, the method is somewhat robust to initial landmark placement.

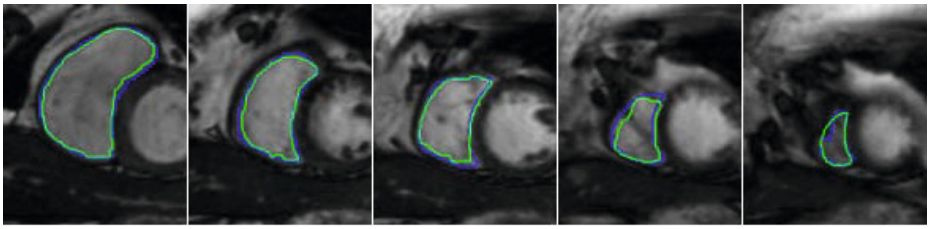

Fig. 3. Example output of our segmentation method. All images are from the same examination at ED. The manual segmentation is outlined in blue and the segmentation produced by the proposed method is outlined in green.

Table 2. Results from RV segmentation and comparison with the state-of-the-art

	Phase	Dice Metric	Haussdorf Distance
Proposed Method	ED	$\mathbf{0.86 \pm 0.10}$	8.40 ± 4.21
	ES	$\mathbf{0.75 \pm 0.18}$	$\mathbf{10.02 \pm 5.78}$
Grosgeorge et al. [10]	ED	0.83 ± 0.15	9.48 ± 5.41
	ES	0.70 ± 0.22	10.56 ± 5.54
Bai's et al. [6]	ED	0.86 ± 0.11	$\mathbf{7.70 \pm 3.75}$
	ES	0.69 ± 0.25	11.16 ± 5.53

For each slice segmented, both the DM and Haussdorf Distance (HD) were computed. The results can be seen in Table 2 where they are compared against Grosgeorge et al. [10] and Bai et al. [6]. Example output can be seen in Fig. 3. The results show that our method performed similarly or slightly better than the state-of-the-art in most metrics, with greatest improvement during ES. This, coupled with the results in Table 1 lends support to our argument that a manifold model is well suited for dealing with the complex RV shapes. The code was written in C++ with little optimisation. Segmentation of a single volume takes between 7-10 minutes depending on the speed of convergence. The tests were run on an Intel Core i7 CPU @ 2.80GHz with 16Gb of memory.

4 Conclusions and Future Work

We proposed a novel method of combining shape priors generated from a learned manifold into a Markov Random Field segmentation framework. By alternating between segmentation and improving the prior, we have shown that the manifold is able to generate priors that fit to the shape of target even when the morphology is complex. We have tested the performance of our method with an independently evaluated dataset with results that are comparable with the state-of-the-art methods.

For future work we intend to implement the method in 3D to be able to use the improved accuracy of the basal segmentations to influence the more difficult

apical slices. In addition, we aim to improve our manifold traversal method by better utilising the underlying geometry of the data. This will allow us to extrapolate, increasing the amount of shape variation that the model can fit to.

References

1. Petitjean, C., Dacher, J.N.: A review of segmentation methods in short axis cardiac MR images. Medical Image Analysis 15(2), 169–184 (2011)
2. Caudron, J., Fares, J., Lefebvre, V., Vivier, P.H., Petitjean, C., Dacher, J.N.: Cardiac MRI assessment of right ventricular function in acquired heart disease: factors of variability. Academic Radiology 19(8), 991–1002 (2012)
3. Petitjean, C., Ruan, S., Grosgeorge, D.: Right Ventricle Segmentation in Cardiac MRI: a MICCAI 2012 Challenge (2012)
4. Warfield, S., Zou, K., Wells, W.: Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. IEEE Trans. Medical Imaging 23(7), 903–921 (2004)
5. Zuluaga, M.A., Cardoso, M.J., Modat, M., Ourselin, S.: Multi-atlas propagation whole heart segmentation from MRI and CTA using a local normalised correlation coefficient criterion. In: Ourselin, S., Rueckert, D., Smith, N. (eds.) FIMH 2013. LNCS, vol. 7945, pp. 174–181. Springer, Heidelberg (2013)
6. Bai, W., Shi, W., O'Regan, D.P., Tong, T., Wang, H., Jamil-Copley, S., Peters, N.S., Rueckert, D.: A probabilistic patch-based label fusion model for multi-atlas segmentation with registration refinement: application to cardiac MR images. IEEE Trans. Medical Imaging 32(7), 1302–1315 (2013)
7. Mitchell, S.: Multistage hybrid active appearance model matching: segmentation of left and right ventricles in cardiac MR images. IEEE Trans. Medical Imaging 20, 415–423 (2001)
8. Ordas, S., Boisrobert, L., Huguet, M., Frangi, A.: Active shape models with invariant optimal features (IOF-ASM) application to cardiac MRI segmentation. Computers in Cardiology, 633–636 (2003)
9. Cootes, T., Taylor, C.: Active shape models-their training and application. CVIU 61, 38–59 (1995)
10. Grosgeorge, D., Petitjean, C., Dacher, J.N., Ruan, S.: Graph cut segmentation with a statistical shape model in cardiac MRI. CVIU 117(9), 1027–1035 (2013)
11. Zhu-Jacquot, J., Zabih, R.: Graph Cuts Segmentation with Statistical Shape Priors for Medical Images. In: Signal-Image Technologies and Internet-Based System, pp. 631–635 (2007)
12. Coifman, R.R., Lafon, S.: Diffusion maps. Applied and Computational Harmonic Analysis 21(1), 5–30 (2006)
13. Etyngier, P., Segonne, F., Keriven, R.: Shape Priors using Manifold Learning Techniques. In: IEEE ICCV, pp. 1–8 (2007)
14. Thorstensen, N., Étyngier, P., Ségonne, F., Keriven, R.: Diffusion maps as a framework for shape modeling. CVIU 115(4), 520–530 (2011)
15. Chan, T., Zhu, W.: Level Set Based Shape Prior Segmentation. IEEE CVPR 2, 1164–1170 (2005)
16. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. IEEE PAMI 23(11), 1222–1239 (2001)
17. Powell, M.: Direct search algorithms for optimization calculations. Acta Numerica (1998)