

Level Set Segmentation of the Fetal Heart

I. Dindoyal^a, T. Lambrou^a, J. Deng^{a,b}, C.F. Ruff^c, A.D. Linney^a, C.H. Rodeck^b,
and A. Todd-Pokropek^a

^aDepartments of Medical Physics and Bioengineering,
^bObstetrics and Gynecology, University College London (UCL),
^cUCL Hospitals NHS Trust (UCLH), UK
{i.dindoyal, t.lambrou, jdeng, cfr, alf,
atoddpok}@medphys.ucl.ac.uk

Abstract. Segmentation of the fetal heart can facilitate the 3D assessment of the cardiac function and structure. Ultrasound acquisition typically results in drop-out artifacts of the chamber walls. This paper presents a level set deformable model to simultaneously segment all four cardiac chambers using region based information. The segmented boundaries are automatically penalized from intersecting at walls with signal dropout. Root mean square errors of the perpendicular distances between the algorithm's delineation and manual tracings are within 7 pixels (<2mm) in 2D and under 3 voxels (<4.5mm) in 3D. The ejection fraction was determined from the 3D dataset. Future work will include further testing on additional datasets and validation on a phantom.

1 Introduction

Congenital heart disease affects about 8 in every 1000 births [1] and its signs can be diagnosed with prenatal echocardiography [2]. As with the adult heart, functional volume estimation of the left ventricle provides quantitative information about the state of the myocardium. However, in the fetus the blood flow in both sides of the heart is allowed to mix and so both ventricles are important for clinical assessment. One important application of fetal cardiac segmentation is for measurement of the absolute size of the chambers. This can be used for evaluation of the function of the heart, compromised either by cardiac malformations or by non-cardiac diseases such as immuno-haemolysis. In this condition the maternal immune system can kill fetal blood cells and so the fetal heart grows larger to compensate.

The prenatal heart has very thin chamber boundaries particularly in the areas consisting of the atrial septum, the membranous segment of the ventricular septum, and the valvular leaflets. Often the resolution of the ultrasound beam perpendicular to its axis is insufficient to resolve these structures and so these walls suffer from signal dropout and appear as holes in the endocardium. These dropouts can also be misleading for clinical diagnosis since the fetal heart contains septal holes which normally close at birth. Artefacts such as these complicate the automated functional volume quantification of each chamber for determining useful cardiac indices such as ejection fraction. In some cases it is difficult for fetal cardiologists to manually trace the endocardiac structures because of the missing image greyscale information.

Automated volume quantification in fetal cardiology is relatively new since it is difficult to acquire datasets without significant shadowing. Fetal body and cardiac motion artefacts are most noticeable when using slice-reconstruction 3D methods. Recent advances in volumetric acquisition have allowed the fetal heart to be imaged in 3D with considerably reduced motion artifacts [3], [4].

In the past Navaux and co-authors have published their work on segmentation of the 2D fetal heart by classification via neural networks [5], [6]. There has been relatively little use of deformable models to segment fetal cardiac data – only two papers in the literature currently exist: Lassige et al [7] used a level set snake to measure the size of the septal defects in echocardiographic images. This snake had a constant speed term that frequently overshoot boundaries. In 2003 Dindoyal and co-workers presented an explicit 2D Gradient Vector Flow (GVF) snake algorithm with rigid body motion constraints to segment and track ventricles in 2D motion-gated fetal cardiac data [8]. Recently Esh-Broder et al [3] have collected over 20 3D fetal heart datasets. In this study the ejection fraction was estimated from the manual segmentation as well as comparison of both left and right ventricular volumes.

Section 2 outlines the proposed method to automatically segment the fetal cardiac chambers of two echocardiographic datasets after placement of manual seed points; one dataset was acquired from conventional 2D ultrasound slices and the other by Live 3D. The next section presents a selection of the images segmented with manual tracings for comparison as well as measurement of the accuracy. We then conclude the work and present further directions for study.

2 Method

2.1 Data Acquisition

3D acquisition of the fetal heart by serial slices was carried out by an online motion gated method pioneered in our group [9] using paired Acuson scanners (25 frames per second and a square pixel spacing of 0.26mm). True 3D acquisition of the fetal heart was performed with the Live 3D ultrasound scanner from Phillips [4]. This imaging system is capable of capturing about 24 volumes per second and can output a resampled cubic voxel resolution of 1.47mm^3 for the penetration depth required. Although the 2D images from the paired scanners were stacked in 3D with motion gating, there was parts of the volume with noticeable motion artifacts that caused misalignment between slices. For this reason the volume dataset from the Acuson scanners was treated as separate 2D images. The Live 3D dataset did not suffer from this problem due to the volumetric acquisition method. Motion gating was unnecessary for this dataset since it was acquired with the probe kept immobile during no apparent fetal body movement.

2.2 Level Set Deformable Model

The level set method is defined implicitly compared to many adaptations of the snake model first introduced by Kass in 1988 [10] which track explicit markers. Level set methods can behave like the explicit case by chopping the level set function at the zero level (refer to Fig. 1). These implicit models have attractive properties in image

segmentation such as automatic interpolation of the propagating front for irregular shaped boundaries and the ability to handle topological changes with ease.




Fig. 1. Illustration of the level set function and zero levels for a 2D image

The generic level set equation for image processing can be written in the form

$$\dot{\phi} = F |\nabla \phi| . \tag{1}$$

where ϕ is the level set function and F is a problem dependent speed function.

The level set used in this paper was solved using a first order iterative scheme because of its low computational complexity

$$\phi_{n+1} = \phi_n + \Delta N F \phi_n . \tag{2}$$

where n is the iteration number and ΔN is the timestep. Upwind differencing schemes were used where appropriate to maintain numerical stability around the propagating front as well as a small timestep. A form of narrow banding was used to speed up the level set propagation and to prevent nucleation of new fronts. The front is tracked on each iteration and its intersection with the edge of the narrow band can be predicted. When this occurs a new band is grown from the current zero level front by isotropic diffusion.

The front was manually initialized as a circle or sphere in each chamber. Each chamber contained a different snake which was stored in separate memory space to the others. This was implemented to prevent the level set merging of neighboring fronts. The distance transform for each snake was defined as a cone with negative values inside the front and positive elsewhere. This can be computed very quickly for such simple geometry in a single pass by computing the distance between voxel positions from the radius of the front. For non primitive initializations it may be necessary to use more general efficient distance transforms such as chamfering. The usual criterion of normalizing the distance transform was enforced.

Sarti et al 2002 [11] developed a level set algorithm with mean curvature and edge flow diffusion properties to segment datasets with missing boundaries. In this paper a new term was added to this evolution equation proposed by Sarti to incorporate region growing based on local deviations from the interior and exterior regions using the image part of the Mumford Shar (MS) functional. This term is useful in images without clear boundaries [12]. In the implementation proposed here the MS force is heavily penalized by curvature and inter-snake collision detection to reduce inter-

chamber leakage. This is shown in equation (3) where Sarti's geometric model for boundary completion is enclosed in curly braces.

$$\begin{aligned} \dot{\phi} = & \left\{ \alpha g \nabla \cdot \left[\frac{\nabla \phi}{|\phi|} \right] + \beta (\nabla g) \cdot (\nabla \phi) \right\} \|\nabla \phi\| \\ & + \left(\lambda_1 [I - \mu_i]^2 - \lambda_2 [I - \mu_o]^2 \right) \exp \left(-\kappa \nabla \cdot \left[\frac{\nabla \phi}{|\phi|} \right] \right) (1 - \xi) \|\nabla \phi\|. \end{aligned} \quad (3)$$

Where $g = (I + |\nabla GF|)^{-1}$ is an edge detector that returns a value between 0 and 1, with G denoting Gaussian filtering and F is the image. In the implementation for this paper image prefiltering to reduce noise was unnecessary for the volumetric data since the images were already at very low spatial resolution, but was necessary for the 2D sliced data due to the high speckle content.

In equation (3) ϕ is the level set function, I is the current voxel intensity under investigation. μ_i and μ_o are the means of the internal and outside regions of the dataset defined by the level set front. ξ is a function that tests if any of the enclosed regions from individual snakes overlap. If there is overlap ξ returns 1 and 0 otherwise. β is a function to penalize edge advection in the presence of local edges and is defined as $\beta = \exp(-\gamma |\nabla GF|)$. In Sarti's original formulation β is a unitary constant. The factors α , γ , λ_1 , λ_2 , κ are empirically determined weighting coefficients for the respective terms.

In Sarti's equation the first term is standard mean curvature flow weighted by an edge stopping coefficient. It serves to regularize the curve where the data is sparse and propagation can be further reduced by the presence of edges. The advection term drives the front towards image edges that have been defined from a pre-computed edge diffusion field. The main weakness of this term is the presence of many edges at various strengths as is often found in sonography. Edge flow by advection is heavily dependent on the quality of the edgemap and so may fail to propagate the front towards the edges sufficiently to overcome the mean curvature flow.

The proposed term aims to provide some expansion or contraction forces dependent on the local tissue type in the absence of a strong edge field, e.g. when the front is in homogeneous regions. Unlike the constant advection term in Lassige's algorithm [7] this force can propagate the front in either direction according to the position of the boundaries and so would be less prone to overshoot. The MS factor models the foreground and background of the image and tries to minimize its energy by separating these two regions. The foreground was estimated from a small circle/sphere placed inside the chamber prior to evolution and the background was assumed to be the remainder of the dataset. Since the appropriate λ_1 and λ_2 could potentially vary significantly between datasets, the images were normalized to reduce the dependence on these coefficients.

The exponential factor contains a second mean curvature component and its presence is mostly required where there is extensive shadowing to the chamber. The collision detection component is heavily penalizing and tends to stop two intersecting fronts immediately upon contact so that a steady wall is formed where the two interfaces meet. Although open valves cause blood from atria and ventricles to mix cardiac function in clinical use is measured by treating each chamber in isolation.

From preliminary experiments it was discovered that for the collision to occur at the right place (where part of the chamber wall has suffered signal dropout due to the beam resolution); the two snakes should be started from as close to the centers of their respective chambers as possible. This prevents one snake from invading the adjacent chamber due to its arrival at the missing boundary first.

3 Results and Discussion

Fig. 2 illustrates the effect of the added term to Sarti's equation. Without the presence of a clear edgemap from the data Sarti's snake fails to propagate appreciably towards the desired boundary. To overcome this restriction we used the bidirectional MS term in conjunction with Sarti's algorithm which yielded a closer segmentation to the expert's delineation. We used manual expert tracings by a fetal cardiologist as a gold standard. Full interactive segmentation of the images proved to be both challenging and tedious in particular for areas with partial volume artifacts of the papillary muscles and the missing atrial septum.




Fig. 2. Segmentation by the algorithm proposed in this paper (*left*) and segmentation by Sarti's algorithm (*middle*). The white contours are automatically generated and grey denotes manual tracings. Atria appear at the top of the image and ventricles at the bottom. The *right* image shows the edgemap




Fig. 3. Effect of the collision penalization term (enabled in *left* image and disabled in *right* image)

Fig. 3 shows the effect of the collision penalization term. The left image shows the atrial boundary reconstructed at approximately the location of the true boundary and the right image illustrates overgrowth of the fronts.




Fig. 4. Segmentation results of the algorithm on the 2D slice data (*white*) superimposed on manual tracings (*grey*)




Fig. 5. 3D segmentation on the Phillips Live 3D data (only central slices are shown for clarity). Space varies *horizontally* and time *vertically* in this figure. *White* contours were generated by the algorithm and *grey* ones are manual tracings

The algorithm was applied to 32 2D images from the Acuson scanner and a 3D dataset from the Phillips Live 3D scanner. Inter-operator variability of manual segmentation was not measured for comparison with the automated delineation in this paper. It was assumed that the repeatability of manual tracing to within an error of 2 voxels would be sufficient for clinical use. To assess the accuracy of the algorithm

root mean square (rms) errors were computed from point-wise distances between the automatic boundaries and the manual tracings. The segmentation results in 2D can be seen in Fig. 4 and 3D analysis is displayed in Fig. 5.

Frequency analysis of the rms errors is shown in the bar charts in Fig. 6. The rms errors are within 7 pixels (<2mm) in 2D for the chambers excluding the right atrium. In Fig. 6 the scale was truncated to empathize the distribution of 2D rms errors excluding outliers of the right atrium. The outliers stretch out to 15 pixels. Fig. 4 shows several examples where the snake is attracted to regions of echo enhancement in the ultrasound image above the right atrium. The algorithm appears to be influenced by strong intensity inhomogeneity artifacts and this pulls the contour away from the desired endocardiac boundaries.

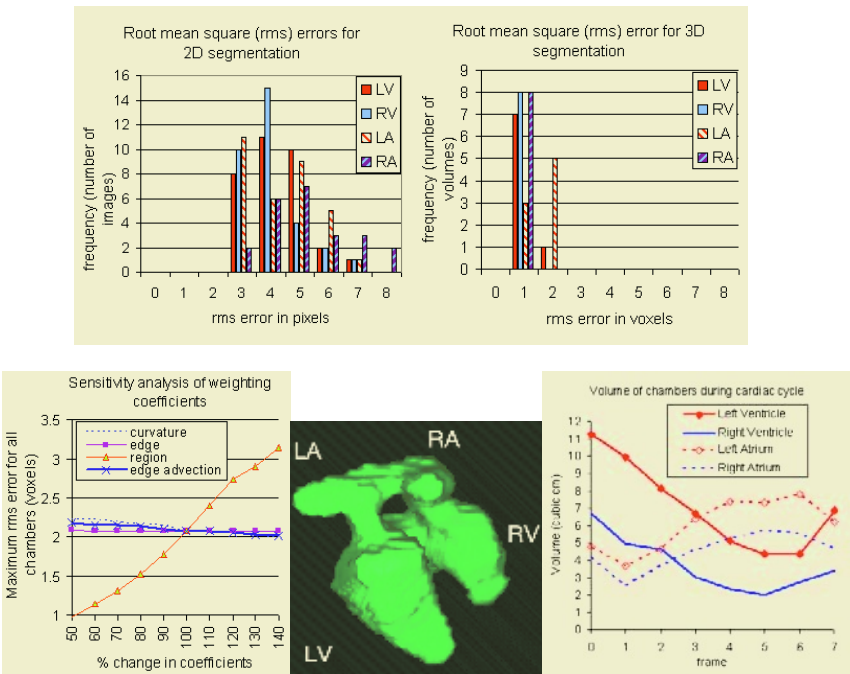


Fig. 6. Top row: rms errors of segmentation for each chamber. Bottom left: sensitivity analysis of the weighting coefficients in the 3D segmentation. Bottom middle: example volume rendering of segmented Live 3D dataset (during ventricular diastole and atrial systole, posterior view). Bottom right: volume-time curves for all frames

In 3D the rms errors were between 1-2 voxels (4.5mm). The graphs show that the spread of rms errors is greater in the 2D case although the overall errors from the 3D algorithm are larger when expressed in mm units. We attribute the cause of this to the lower spatial resolution in direct volumetric acquisition which increases the partial

volume effect. Strong echo enhancement effects would further increase the spatial rms errors but these were absent in this dataset.

In many of the 2D segmented images the front stopped short of the desired boundary and this shows up in the high rms errors (refer to Fig. 6). Whilst high curvature penalization was partially responsible, premature stopping of the level set front was also due to contributions from the type of image forces used. The MS term models unchanging mean intensities both inside and outside of a small sphere or circular seed placed inside the chamber. If the mean was updated as the snake evolved the front could come to rest closer to the boundary. The edge flow term requires a diffusion equation to be applied to an edge map and so broadens edges. The Gaussian prefiltering used a large kernel of width 9 pixels and this could contribute to the edge broadening.

To assess the dependence of the segmentation on the weighting coefficients, these parameters were varied independently within a range of 50-140% of their appropriate values for this application. The test was arbitrarily run on the 1st frame of the Live 3D dataset and for this range of parameter values, rms errors in the segmentation remained below 3.5 voxels for all chambers. Fig. 6. shows that the MS region term is most strongly affected by choice of weighting coefficients.

The left and right ventricular ejection fractions were computed as 61% and 59% respectively from the automatic segmentation (frames 0 and 6 were identified as the necessary cardiac time points to perform this calculation). These values are comparable to the ranges found in a study by Esh-Broder [3] ($57.5\pm 14.6\%$ and $54\pm 11.2\%$ respectively), who also measured a non-significant variation between the left and right ventricles. A volume-time graph of all four cardiac chambers is shown in Fig. 6 as well as a volume rendering of the level set front.

4 Conclusion

To our knowledge this is the first time non manual segmentation techniques have been applied on 3D prenatal heart data to measure volumes and cardiac indices. The automated method provides a segmentation that is far quicker and more repeatable than manual tracings; but has problems in delineating fine intra-cavity structures and is strongly affected by enhancement of echogenic regions. The 2D data appears to be easier to automatically segment than the new Live 3D possibly due to the smaller pixel size which allows better definition of thin walls and intra-cavity structures. However, fetal cardiac volumetric imaging simplifies the gating process and will be the method of choice for acquiring future datasets.

Future work will involve validation of the accuracy of the automated volume determination from a phantom as well as further testing on additional datasets.

Acknowledgements

We are grateful for the volume rendering provided by Daren McDonald and stimulating discussions about optimization strategies with Dr Robin Richards. This work was supported by EPSRC (GR/N14248/01) and MRC (D2025/31) under the

Interdisciplinary Research Consortium scheme - "From Medical Images and Signals to Clinical Information" (MIAS IRC). Dr Jing Deng is supported by MRC (G108/516).

References

1. Mitchell SC, Korones SB, Berendes HW. Congenital heart disease in 56,109 births. Incidence and natural history. *Circulation*, Vol. 43. (1971) 323-332
2. Copel JA, Gianluigi P, Green J, Hobbins JC, Kleinman CS. Fetal echocardiographic screening for congenital heart disease: The importance of the four-chamber view. *British Journal of Obstetrics and Gynaecology*, Vol. 157. (1987) 648-655
3. Esh-Broder E, Ushakov FB, Imbar T, Yagel S. Application of free-hand three-dimensional echocardiography in the evaluation of fetal cardiac ejection fraction: a preliminary study. *Ultrasound in Obstetrics & Gynecology*, Vol. 23. (2004) 546-551
4. Deng J. Terminology of three-dimensional and four-dimensional ultrasound imaging of the fetal heart and other moving parts. *Ultrasound in Obstetrics & Gynecology*, Vol. 22. (2003) 336-334
5. Piccoli L, Dahmer A, Scharcanski J, Navaux POA. Fetal echocardiographic image segmentation using neural networks. In: *Image Processing and its Applications 1999, Seventh International Conference*. (1999)
6. Siqueira ML, Scharcanski J, Navaux POA. Echocardiographic image sequence segmentation and analysis using self-organizing maps. *Journal of VLSI Signal Processing*, Vol. 32. (2002) 135-145
7. Lassige TA, Benkeser PJ, Fyfe D, Sharma S. Comparison of septal defects in 2D and 3D echocardiography using active contour models. *Computerized Medical Imaging and Graphics*, Vol. 24. (2000) 377-388
8. Dindoyal I, Lambrou T, Deng J, Ruff CF, Linney AD, Todd-Pokropek A. An active contour model to segment foetal cardiac ultrasound data. In: *Medical Image Understanding and Analysis*. 03 Jul 10, University of Sheffield, UK. (2003)
9. Deng J, Ruff CF, Linney AD, Lees WR, Hanson MA, Rodeck CH. Simultaneous use of two ultrasound scanners for motion-gated three-dimensional fetal echocardiography. *Ultrasound in Medicine and Biology*, Vol. 26. (2000) 1021-1032
10. Kass M, Witkin A, Terzopoulos D. Snakes: Active Contour Models. *International Journal of Computer Vision*, Vol. 1. (1988) 321-331
11. Sarti A. Subjective surfaces: a geometric model for boundary completion. *International Journal of Computer Vision*, Vol. 46. (2002) 201-221
12. Gibou, F. and Fedkiw, R. A fast hybrid k-means level set algorithm for segmentation. *Stanford Technical Report*. (2002)