Localizing Cardiac Structures in Fetal Heart Ultrasound Video

Christopher P. Bridge[a], Christos Ioannou[b], and J. Alison Noble[a]

[a]Department of Engineering Science, University of Oxford, [b]John Radcliffe Hospital, Oxford

Summary

- We build on our previous particle filtering algorithm for automatic analysis of ultrasound videos from fetal heart screening scans.
- We extend the previous method to track multiple anatomical structures of interest.
- *Partitioning* the state of the particle filter results in more efficient use of particles in the resulting high-dimensional state space.
- Fourier model captures near-periodic motion due to the cardiac cycle.
- Resulting algorithm can track multiple structures in the heart through motion and view changes in a highly efficient manner.

Background

In previous work [1] we addressed the problem of automatic analysis of fetal cardiac screening videos as a first step towards improving the detection rates of congenital heart disease. In that work, we showed how a model based on particle filtering could estimate important 'global' information about each frame: the visibility h, location $\mathbf{x}$ and orientation θ of the heart, the viewing plane classification v, and the position in the cardiac cycle ϕ in an efficient and fully automated manner.

However, many forms of CHD affect specific parts of the cardiac anatomy, such as a valve or vessel. Therefore in this work, we present a direct extension to the previously published model to allow it to track the locations of important structures of the cardiac anatomy, shown in Figure 1, to allow them to be inspected for anomalies.

Structures of Interest

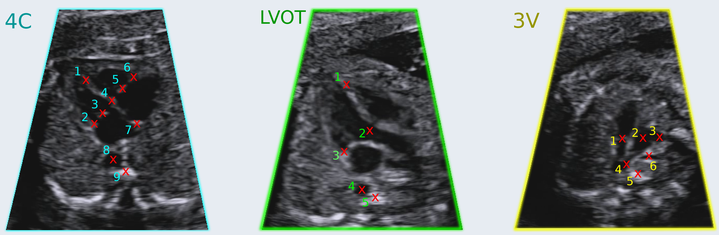

4C: 1. *Apex*, 2. *Mitral Valve End*, 3. *Mitral Valve Centre*, 4. *Crux Cordis*, 5. *Tricuspid Valve Centre*, 6. *Tricuspid Valve End*, 7. *Base*, 8. *Descending Aorta*, 9. *Spine*

LVOT: 1. *Apex*, 2. *Aortic Valve*, 3. *Mitral Valve End*, 4. *Descending Aorta*, 5. *Spine*

3V: 1. *Pulmonary Valve*, 2. *Ascending Aorta*, 3. *Superior Vena Cava*, 4. *Descending Aorta*, 5. *Spine*, 6. *Trachea*

Figure 1: Structures of interest and the three viewing planes used in experiments.

Filter for 'Global' Variables

The filter in [1] tracked heart state $\mathbf{s}_t = (h_t, v_t, \mathbf{x}_t, \theta_t, \phi_t)$ at time t given image information, $\mathbf{z}_t$. The key components are:

- **Prediction Potential:** $\psi(\mathbf{s}_t, \mathbf{s}_{t0-1})$ measures probability of the current state given the previous state.
  - Built using standard probability distributions.
- **Observation Potential:** $\omega(\mathbf{s}_t, \mathbf{z}_t)$ measures compatibility of the current state estimate and image information.
  - Built using random forest classifiers/regressors and rotation invariant features.

If the filtering distribution over the state is approximated by a finite number of weighted samples ('particles'), then updating the particle set at each time step involves applying three operations in sequence:

- **Convolution** ('$*$') with the prediction potential
- **Re-weighting** ('$\times$') with the observation potential
- **Re-sampling** ('$\sim$') the particle set according to the particle weights

We can display these operations in a filtering diagram [2] like so:

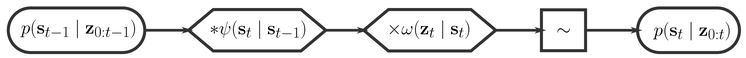

Figure 2: Filtering diagram for a basic particle filter.

Partitioned Particle Filter

To track structures to this architecture, we could just add their locations to the existing state space. However, particle filters do not perform well in high-dimensional state spaces, because an increasingly large number of particles are needed. To address this, we first *partition* the state space of the existing model (making certain assumptions) to give three partitions that can be operated on in sequence [2]. This guides particles through the peaks of the distribution during the resampling process, and therefore makes *more efficient use of a small number of particles.*

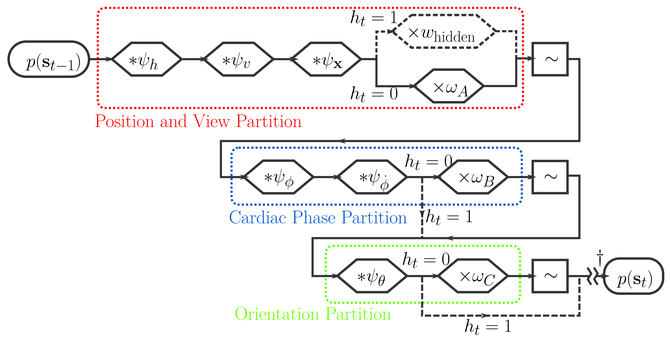

Figure 3: Filtering diagram for the partitioned archtitecture for tracking 'global' variables.

Here, there are independent prediction potentials (ψ_h, ψ_v, etc.) for each state variable and three observation potentials: a classification/detection forest (ω_A), a cardiac phase regression forest (ω_B), and an orientation regression model (ω_C). Hidden particles are given a fixed weight w_{hidden}, and avoid later re-weighting stages.

Fourier Model for Cardiac Cycle Modelling

To model the near-periodic motion of the structures due to the cardiac cycle, we employ a simple Fourier series model of the position $\mathbf{q}_{a,t}$ of structure a at time t relative to the heart centre $\mathbf{x}_t$:

$$\mathbf{p}_{a,t} = \begin{bmatrix} c_{a,1,1} & c_{a,2,1} \\ c_{a,1,2} & c_{a,2,2} \\ c_{a,1,3} & c_{a,2,3} \\ c_{a,1,4} & c_{a,2,4} \\ c_{a,1,5} & c_{a,2,5} \\ \vdots & \vdots \end{bmatrix}^T \cdot \begin{bmatrix} 1 \\ \cos\phi_t \\ \sin\phi_t \\ \cos 2\phi_t \\ \sin 2\phi_t \\ \vdots \end{bmatrix}$$

$$= \begin{bmatrix} \mathbf{c}_{a,1} & \mathbf{c}_{a,2} \end{bmatrix}^T \cdot \boldsymbol{\phi}_t$$

The coefficient vectors $\mathbf{c}_{a,1}$ and $\mathbf{c}_{a,2}$ are added to the state of the filter and allowed to vary smoothly over time according to a prediction potential $\psi_{\mathbf{c}_a}$. A normal distribution for the parameter vectors are fitted during training. The absolute positions can be recovered using the heart position $\mathbf{x}_t$, radius r and rotation matrix $\mathbf{R}_{[\theta_t]}$:

$$\mathbf{q}_{a,t} = r\mathbf{R}_{[\theta_t]}\mathbf{p}_{a,t} + \mathbf{x}_t$$

Filter for Structure Tracking

A new partition is created for each structure, including prediction potentials for the combined coefficients $\tilde{\mathbf{c}}_a$ and the visibility of the structure g_a, and an observation potentials ω_D are based on a classication forest that shares features with the other forests.

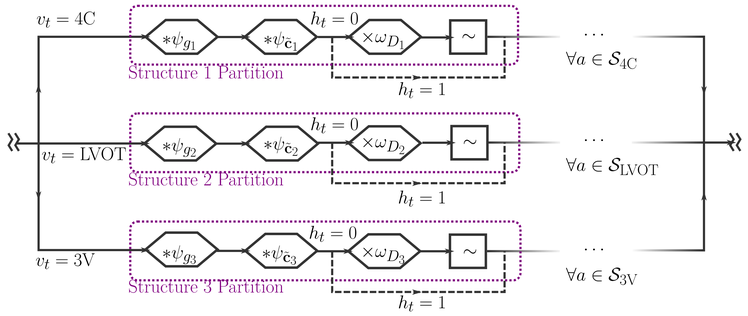

Figure 4: Extension to the filtering architecture for tracking anatomical structures. This extension is inserted at the '$\dagger$' icon in Figure 3.

Experiments

- Database of 91 videos from 12 subjects with manual annotations.
- Leave-one-subject-out cross-validation.
- Multiple views and range of gestational ages (20–35 weeks), orientations, and magnifications.
- Evaluation metric is the localisation error (between detected an ground truth position), normalised by the heart radius.

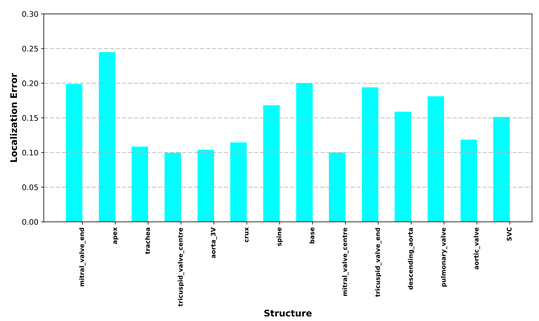

Figure 5: Localisation error, normalised by heart radius, for each structure.

- Many structures are well localised, but some poorly-defined structures show large variation.

Acknowledgements

Christopher Bridge is funded by an EPSRC Doctoral Training Award.

References

[1] C. P. Bridge, C. Ioannou, and J. A. Noble. "Automated Annotation and Quantitative Description of Ultrasound Videos of the Fetal Heart". In: *Medical Image Analysis* 36 (Feb. 2017), pp. 147–161.

[2] J. MacCormick and M. Isard. "Partitioned Sampling, Articulated Objects, and Interface-Quality Hand Tracking." In: *ECCV (2)*. Ed. by D. Vernon. Vol. 1843. Lecture Notes in Computer Science. Springer, 2000, pp. 3–19.