

GRID-ENABLED AUTOMATIC CONSTRUCTION OF A TWO-CHAMBER CARDIAC PDM FROM A LARGE DATABASE OF DYNAMIC 3D SHAPES

S. Ordas¹, L. Boisrobert¹, M. Bossa¹, M. Huguet², M. Laucelli³, S. Olmos¹, AF. Frangi¹

¹Computer Vision Group, Aragon Institute of Engineering Research, University of Zaragoza, Spain.

²Centre Cardiovascular, CETIR Sant Jordi, Barcelona, Spain.

³GridSystems S.A., Palma de Mallorca, Spain.

ABSTRACT

Point Distribution Modelling (PDM) is an efficient generative technique that can be used to incorporate statistical shape priors into image analysis methods like Active Shape Models (ASMs) or Active Appearance Models (AAMs). They are described by a set of landmarks usually manually pinpointed in a training set. Frangi *et al.* [1] have proposed an automatic auto-landmarking technique capable of dealing with multi-object arrangements. In this paper, we present an experimental extension of this previous work, validating the method provided. Our contributions can be summarized as follows: A two-chamber shape model of the heart is constructed from a large data-set comprising 90 subjects and considering 5 phases of the cardiac cycle. The computational demand of our technique is addressed using Grid computing. The results of our experiments suggest that the method presented in [1] as a proof-of-concept, can truly cope with the large inter-subject and inter-phase deformations present in clinical cardiac data sets including pathologies. The achieved accuracy in our validation is comparable to the former tests.

1. INTRODUCTION

Magnetic Resonance Imaging (MRI) is a promising modality for one-stop-shop cardiac examination thanks to its increased spatial and temporal resolution and its ability to provide quantitative morphological and functional information of the heart. An inevitable step before pursuing any kind of quantitative and/or functional analysis, is the segmentation of the cardiac chambers. As the amount of data in dynamic 3D cardiac scans is very large, manual segmentation is not viable, and thus automatic methods are required. During the last few years, model-driven methods and, in particular, statistical 3D models, are being developed for 3D cardiac image segmentation [2], [3]. In a cardiac 3D Point Distribution Model (PDM), a set of landmarks is positioned in the endo and epi boundaries of the ventricles. These landmarks have to be placed in a consistent way over a large database of training shapes to ensure that the final model gathers representative statistics of the shape population. Manual landmarking of dynamic 3D structures like the heart is basically unfeasible due to the large number of landmarks and training shapes that are required

This work is funded by MAPFRE Medicina Foundation, and partially supported by TIC2002-04495-C02, ISCIII G03/185, and FIT-070000-2003-585 grants. SeO and LB hold a MEC-FPU grant AP2002-3955 and AP2002-3946 respectively. AFF holds a MCyT Ramon Y Cajal Fellowship. The authors acknowledge the technical support team of GridSystems, for their help with the InnerGrid Nitya Middleware.

for the construction of detailed spatio-temporal models. The purpose of this paper is thus to further validate the approach in [1] and to investigate the properties of a 3D+t bi-ventricular model constructed from a large population of healthy and diseased hearts, at different instants of the cardiac cycle.

2. THEORY

2.1. Statistical Shape Models

Consider a set $\mathcal{X} = \{\mathbf{x}_i; i = 1 \dots n\}$ of n shapes. Each shape is described by the concatenation of m 3-D landmarks $\mathbf{p}_j = (p_{1j}, p_{2j}, p_{3j}); j = 1 \dots m$, obtained from a surface triangulation. \mathcal{X} is thus a distribution in a $3m$ -dimensional space. The goal is to obtain a general and reasonably compact representation of the population, learnt from the training set. This representation allows to approximate any shape using the following linear model

$$\mathbf{x} = \hat{\mathbf{x}} + \Phi \mathbf{b} \quad (1)$$

where $\hat{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i$ is the average landmark vector, \mathbf{b} is the shape parameter vector of the model, and Φ is a matrix whose columns are the principal components of the covariance matrix $\mathbf{S} = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{x}_i - \hat{\mathbf{x}})(\mathbf{x}_i - \hat{\mathbf{x}})^T$. The principal components of \mathbf{S} are calculated as its eigenvectors, ϕ_i , with corresponding eigenvalues, λ_i (sorted so that $\lambda_i \geq \lambda_{i+1}$). If Φ contains only the first $t < \min\{m, n\}$ eigenvectors corresponding to the largest non-zero eigenvalues, we can approximate any shape of the training set, \mathbf{x} , using Eq. (1) where $\Phi = (\phi_1 | \phi_2 | \dots | \phi_t)$ and \mathbf{b} is a t dimensional vector given by $\mathbf{b} = \Phi^T (\mathbf{x} - \hat{\mathbf{x}})$. Assuming that the cloud of landmark vectors follows a multi-dimensional Gaussian distribution, the variance of the i -th parameter, b_i , across the training set is given by λ_i . By varying these parameters, different instances of the shape class under analysis can be generated using Eq. (1). By applying limits to the variation of b_i , usually $|b_i| \leq \pm 3\sqrt{\lambda_i}$, it can be enforced that a generated shape is similar to the shapes contained in the training class.

2.2. Automatic Landmarking Method

The general layout of the method is to align all the shapes of the training set to an atlas that can be interpreted as a mean shape. Once all the necessary transformations are obtained, they are inverted and used to propagate any number of arbitrarily sampled landmarks on the atlas, to the coordinate system of each subject. In this way, while it is still necessary to manually draw the contours in each training image, our technique relieves from manual landmark

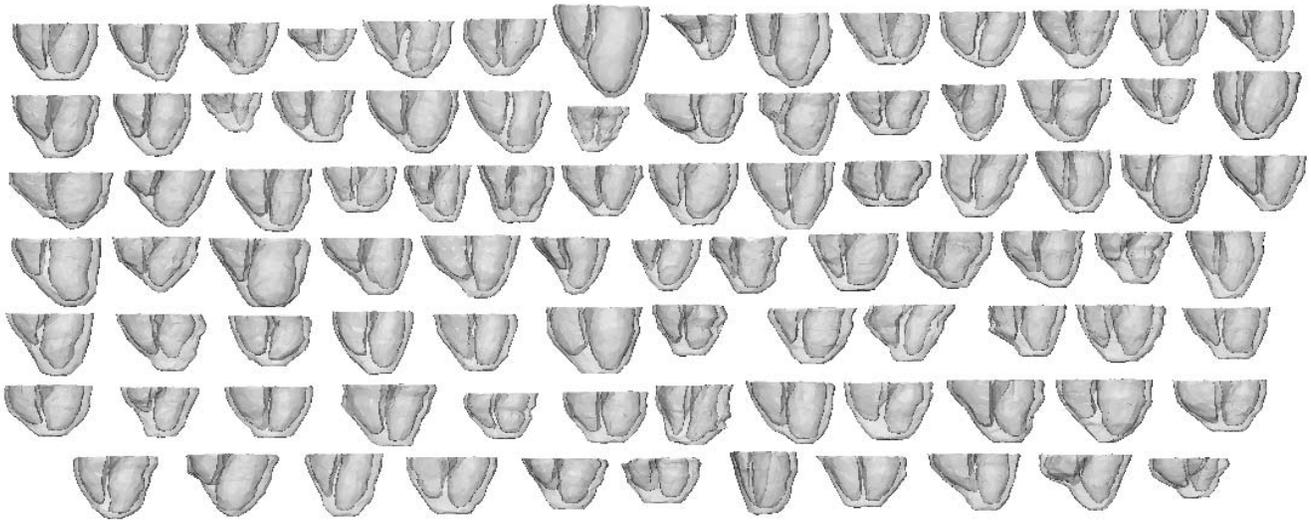


Fig. 1. A snapshot of the database (90 shapes) at ED.

definition and for establishing the point correspondence across the training set. A detailed description of the method can be found in [1], but can be summarized as follows :

1. Manual segmentations of the training set are shape-based interpolated to obtain isotropic voxels.
2. The resulting shapes (Fig. 2) are aligned through a quasi-affine registration adapted from [4] (nine degrees of freedom: translation, rotation, and anisotropic scaling) to a Reference Sample (RS). This sample is randomly chosen from the training set and is considered the first atlas estimate.
3. A new atlas is constructed by shape-based averaging of the aligned shapes and taking the zero-isosurface.
4. To minimize the bias introduced by the choice of the RS, steps 2 and 3 are repeated until the atlas becomes stable. At this point, the atlas is said to be in a Reference Coordinate System (RCS).
5. Subsequently, each aligned sample shape is non-rigidly registered to the RCS atlas.
6. The obtained transformations are then averaged and the resulting averaged transformation is applied to the atlas in RCS. The new atlas is said to be in a Natural Coordinate System (NCS) (see Fig. 3 for a visual comparison). The NCS atlas is unique regardless the RS election [1].
7. A new set of quasi-affine and non-rigid transformations are recalculated in the same way as in 2 and 5.
8. Finally, any automatically generated landmarks in the NCS atlas can be propagated to the samples shapes through the transformations in 7.
9. In order to build the statistical shape model, a Procrustes shape alignment and a Principal Component Analysis (PCA) of the transformed landmarks is performed.

In all the registrations of the method (i.e. of multi-valued images), Label Consistency (LC) [1] is used as the similarity



Fig. 2. Shape-based interpolated images

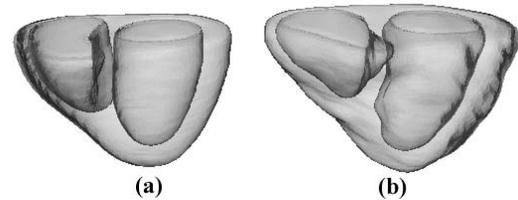


Fig. 3. RCS (a) and NCS (b) atlases

metric. For the non-rigid registrations, a multi-resolution Free Form Deformation (FFD) registration algorithm is used [5], [6]. After the mesh is deformed at each resolution level, it can be refined using a B-spline basis subdivision scheme. The number of levels and initial resolution of the FFD is chosen as a compromise between computational complexity (increasing with the number of resolutions) and accuracy (improving with a smaller mesh spacing) of the deformation. In our application, an initial resolution of $16 \times 16 \times 20 \text{mm}^3$ and 4 resolution levels were used.

3. MATERIALS

3.1. Data Set Description

Our data set (see Fig. 1) contains 90 MR studies from CETIR Sant Jordi Centre (Barcelona, Spain) randomly selected from those acquired during 2002. These studies correspond to 21 healthy subjects, and 74 patients suffering from common cardiac pathologies including myocardium infarction (25), hypertrophy (21), LV dilation (6), LV aneurysm (2), RV dilation (2), LA dilation (5), RA dilation (2), and pericarditis (4), among others. The type, number, and relative proportion of pathologies, is representative of typical examinations in the health center. For each study, only the short axis view was considered. Expert segmentations were manually drawn on three closed contours, namely the endocardial left ventricle (LV ENDO) and right ventricle (RV ENDO) borders, and the epicardial border of the whole heart (HEART EPI). The 2-chamber model included 8-12 slices from the base to the apex. Following common clinical practice, LV ENDO and RV ENDO contours did not include the papillary muscles and trabeculae, and the HEART EPI contour was drawn along the inner border of the epicardial fat layer. The base slice was defined as the most basal slice with the RV still present. The acquisition parameters were: TR: 3.75 4ms, TE: 1.5-1.58 ms, FA: 45, slice thickness: 8-10 mm, slice size: 256×256 pixels, resolution: 1.56×1.56 mm and FOV: 400×300 mm², on a General Electric CVI 1.5 T MR facility. Five phases of the cardiac cycle were segmented. They are denoted : (i) ED (End Diastole), (ii) MS (Mid Systole), (iii) ES (End Systole), (iv) D1 (Diastole 1), and (v) D2 (Diastole 2).

3.2. Grid Computing

A well-known limitation of any non-rigid registration approach is its high computational load. A single registration process (a rigid registration, followed by a non-rigid elastic registration) takes more than two hours in modern CPUs. When dealing with large databases this issue becomes a serious problem for a wide-spread applicability of the method. Distributed computing constitutes a feasible solution. Among several possibilities, Grid computing provides an easy way of taking benefit from off-the-shelf computer clusters. Our Grid middleware platform is the InnerGrid Nitya developed by GridSystems, running on a 13-node dual Xeon (2.8 GHz CPU, 2 GHz RAM) cluster, under Linux RedHat 9. In combination, the cluster represents more than 36 GFlops. With the described facility, for instance, the construction of an atlas of all the phases (450 shapes) with 5 iterations, took approximately 3 hours. Considering that the average processing time for a rigid registration of one sample to an atlas is 4.5 minutes and that the shape-blending, pre- and post- processes tasks takes 16 minutes, this time would represent one week of calculation for a single CPU. For the non-rigid registration of the same 450 shapes, the gridified process lasted 56.5 hours, while for a single CPU it would have taken 1483 hours (two months).

4. EXPERIMENTS

The idea behind the design of the experiments was two fold. On the one hand, it was interesting to test if the method could cope with large data sets with clinically representative inter-subject and inter-phase variability. On the other hand, the idea was to explore the statistical value of our model by testing to which extent the

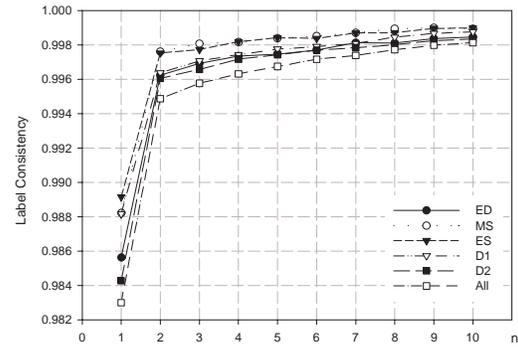


Fig. 4. Convergence of the atlas construction algorithm - evolution of the LC similarity metric over iterations.

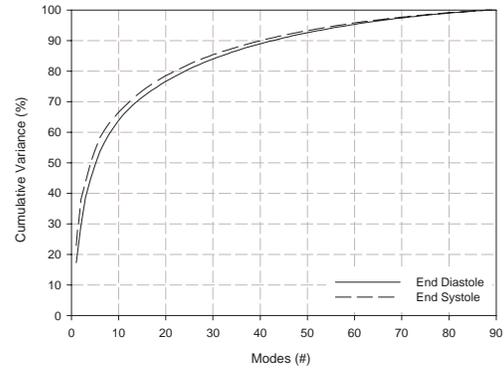


Fig. 5. Percentage of total shape variance versus the number of modes used in the ES and ED shape models.

number of patients in the training set was capable of describing the shape variation of the healthy/diseased heart population.

Six different atlases were constructed. One for each phase (5 atlases with 90 shapes each) and another atlas (denoted ALL) including all the shapes of the five phases (450 shapes). The iterative process adopted to construct the atlas has shown to converge rapidly. Fig. 4 shows that two iterations were enough to give a very high LC. The behavior of all the atlases was similar and tended to stabilize at the same value. The atlas of the fifth iteration was used for the following stages of the process. The influence of the mesh density was also studied. Five mesh densities with 9603, 4998, 2583, 1333, and 694 landmarks were produced. The creation of surface models with correct topology and optimized triangular shape from the binary atlas volume was carried out automatically with an improved marching cubes algorithm. Adaptive mesh decimation was performed to achieve meshes with less number of elements. Fig. 5 shows that the variance explained for the ED and ES models is similar. For instance, to explain 98% of shape variance, 73 and 74 modes are necessary in the ED and ES models, respectively. The curves for different number of landmarks (not shown) are almost coincident. A point to emphasize from here is that regardless the mesh density (as long as the topology is reasonably preserved), the shape variation is explained in the same way. The reconstruction error of the shape model was assessed by means of a leave-one-out test. Fig. 6 shows the results of the

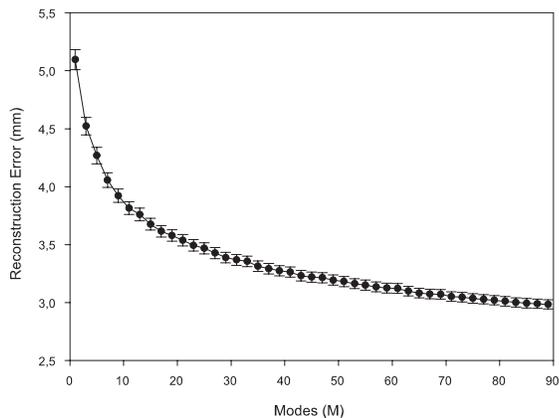


Fig. 6. Reconstruction error in a leave-one-out experiment.

experiment carried out for the ED model with 2110 landmarks. The behavior of other phases and mesh densities was very similar. The reconstruction error, although not very high, does not stabilize with larger number of modes. This may tell us that 90 subjects is not yet large enough to appropriately model the shape variation in the corresponding cardiac phase. Therefore, a larger training set can still improve the generalization ability of the model.

In Fig. 7 the principal modes of variation of the ED model are illustrated. The visual behavior of their variability is quite similar for different mesh densities. It is also worth to emphasize that the landmarks are propagated smoothly, preserving the topology, without triangle flipping or surface folding.

5. CONCLUSIONS

The presented work shows that the method presented in [1] as a proof-of-concept, can truly cope with large heterogeneous data sets of cardiac MR studies. A compact middleware Grid application was designed to deal with the large computational cost of the method. The generalization capability of the approach, suggests feasibility for the use of other modalities (e.g. SPECT, CT) and organs with shape variability not differing too much from that of the heart (e.g. liver, kidneys). Future efforts will go in the line of adding more subjects in the training set and other image planes (e.g. long axis views) to the model. In this way, the shape model would benefit from shape information with increased resolution, yielding a more compact atlas representation. Finally, saying that the statistical shape models are intended to be included in a fitting strategy capable of exploiting the achieved landmark correspondence.

6. REFERENCES

- [1] A.F. Frangi, D. Rueckert, J.A. Schnabel, and W.J. Niessen, "Automatic construction of multiple-object three-dimensional statistical shape models: Application to cardiac modeling.," *IEEE Trans Med Imaging*, vol. 21, no. 9, pp. 1151–1166, 2002.
- [2] S.C. Mitchell, J.G. Bosch, B.P.F. Lelieveldt, R.J. van der Geest, J.H.C. Reiber, and M. Sonka, "3D active appearance

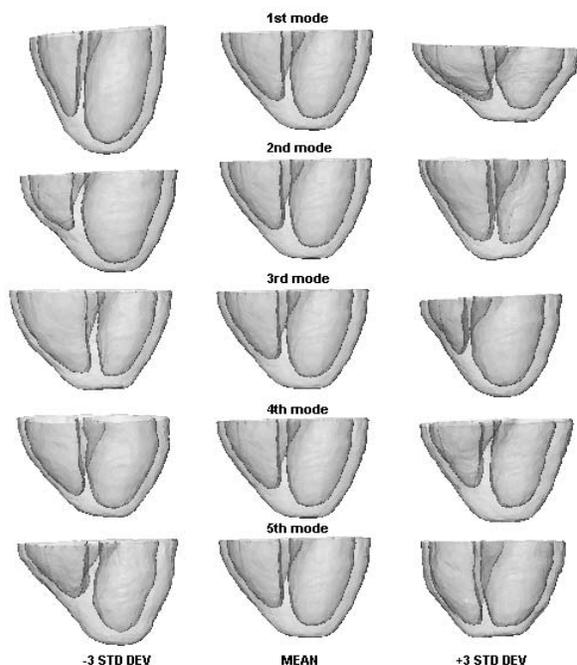


Fig. 7. The five principal modes of variation of the ED model. Each shape instance is generated by varying a single shape parameter while setting the others to zero standard deviations from the mean shape.

- models: Segmentation of cardiac mr and ultrasound images.," *IEEE Trans Med Imaging*, vol. 21, no. 9, pp. 1167–1179, 2002.
- [3] H.C. van Assen, M.G. Danilouchkine, F. Behloul, H.J. Lamb, R.J. van der Geest, J.H.C. Reiber, and B.P.F. Lelieveldt, "Cardiac LV segmentation using a 3D active shape model driven by fuzzy inference," Montreal, CA, Nov. 2003, vol. 2876 of *Lect Notes Comp Science*, pp. 533–540, Springer Verlag.
 - [4] C. Studholme, D.L.G. Hill, and D.J. Hawkes, "Automated 3D registration of MR and PET brain images by multiresolution optimization of voxel similarity measures," *Med Phys*, vol. 24, no. 1, pp. 25–35, 1997.
 - [5] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, "Non-rigid registration using free-form deformations: Application to breast MR images," *IEEE Trans Med Imaging*, vol. 18, no. 8, pp. 712–721, Aug. 1999.
 - [6] J.A. Schnabel, C. Tanner, A.D. Castellano, M.O. Leach, C. Hayes, A. Degenhard, H. Rodney, D.L.G. Hill, and D.J. Hawkes, "Validation of non-rigid registration using finite element methods," Davies, CA, USA, June 2001, vol. 2082 of *Lect Notes Comp Science*, pp. 344–357, Springer Verlag.