

Spatio-Temporal Dimension Reduction of Cardiac Motion for Group-Wise Analysis and Statistical Testing

K. McLeod¹, C. Seiler^{1,2}, M. Sermesant¹, and X. Pennec¹

¹ INRIA Méditerranée, ASCLEPIOS Project, Sophia Antipolis, France

² Institute for Surgical Technology and Biomechanics, University of Bern, Switzerland

Abstract. Given the observed abnormal motion dynamics of patients with heart conditions, quantifying cardiac motion in both normal and pathological cases can provide useful insights for therapy planning. In order to be able to analyse the motion over multiple subjects in a robust manner, it is desirable to represent the motion by a low number of parameters. We propose a reduced order cardiac motion model, reduced in space through a polyaffine model, and reduced in time by statistical model order reduction. The method is applied to a data-set of synthetic cases with known ground truth to validate the accuracy of the left ventricular motion tracking, and to validate a patient-specific reduced-order motion model. Population-based statistics are computed on a set of 15 healthy volunteers to obtain separate spatial and temporal bases. Results demonstrate that the reduced model can efficiently detect abnormal motion patterns and even allowed to retrospectively reveal abnormal unnoticed motion within the control subjects.

1 Introduction

With the continuing high incidence of heart disease around the world in all sectors of society, there is a growing interest in cardiac modeling to aid in diagnosis, prognosis, and therapy planning. Models are continually being developed to try to simulate the different dynamics of cardiac function. Cardiac motion tracking, for instance, is becoming more and more widely studied to try to quantify the different motion dynamics observed between normal and pathological cases.

Over the past few years, a number of methods have been proposed to track the motion of the tissue for different cardiac imaging modalities (ultrasound, cine-magnetic resonance (MR), tagged-MR etc.), with increasingly interesting results (see [1] and references therein). The majority of these methods focus on cardiac motion analysis of a single subject, however, population-based motion analysis may provide further information to aid clinicians. Ideally, such analyses should be robust and reproducible for different subjects drawn from different populations (such as healthy vs. pathological subjects), which requires the motion to be described by few parameters. In [2], a statistical motion model is proposed using

principal component analysis on motion fields to define a per phase and per cycle model. However, the models are used for only individual cases. A statistical atlas of motion computed from healthy volunteers was proposed by Duchateau et. al [3] to quantitatively analyse abnormal motion patterns and aid with therapy planning. This method was effective in characterising abnormal motion, but is limited to 2D ultrasound data. De Craene et. al later extended the method to 4D, but applied the method only to tagged MRI [4]. Recently, a low-dimensional cardiac specific model was proposed by Mcleod et. al [5], that computes the polyaffine fusion over regions estimated from the log-demons deformation with added incompressibility penalisation and neighbouring region regularisation, but again this method is limited to individual cases. Though this method was able to drastically reduce the number of parameters needed to represent the transformation, we would like to go even further to reduce the number of parameters from several subjects to enable robust population-based statistics on the motion.

To that end, we propose to extend on the method of [5] by applying model reduction on the polyaffine model parameters of a single subject (resp. group of subjects) to perform patient (resp. group) -wise statistics on left ventricular motion. The method of [5] is described briefly in the following section for review. Sec. 3 describes a validation of the registration on synthetic sequences with known ground truth, as well as patient-based statistics on the computed affine parameters. In Sec. 4 a method for performing group-wise statistics both spatially and temporally from the Polyaffine model is described and performed on 15 healthy volunteers. Decomposing the motion this way is essentially estimating a tensor decomposition of the spatial and temporal components for multiple subjects, but in a more simplified manner than tensor decomposition (which is not at all straightforward to compute). Similar decomposition was described in [6] and [7], though these required point annotations and temporal correspondences of these points along the cycle. Enforcement of such correspondences can introduce bias since determining such correspondences is not evident. The proposed method requires no such parameterisation and rather using the images directly. The method was able to compute a patient-specific model of the motion with a reduced number of parameters, as well as building separate spatial and temporal bases at a group-wise level to understand the different aspects for a given population and to even detect abnormal motion within that population.

2 Cardiac Motion Tracking with a Near-Incompressible Log-Domain Polyaffine Model

2.1 Cardiac-Specific Polyaffine LogDemos

In Seiler et.al [8], the authors propose a method to estimate a polyaffine model from a log-demons stationary velocity field. Using homogeneous co-ordinates, the parameters of the polyaffine model can be defined for points x in Cartesian co-ordinates as $\log(T) \stackrel{\text{def}}{=} \log \begin{pmatrix} A & t \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} M \\ 0 \end{pmatrix}$, where $\log$ is a principal matrix

logarithm, A is the linear part of the transformation, t its translation, and M a 3×4 matrix. For each segment the affine deformation fields parameterised by the M_i matrices are fused to a global deformation field using the polyaffine model:

$$\mathbf{v}_{poly}(x) = \sum_i \omega_i(x) M_i x, \quad (1)$$

where ω_i is a parameter controlling the weight of the i^{th} region for each voxel x . Eqn. 1 can be estimated by projecting a given observed velocity field $\mathbf{v}(x)$ onto the space of polyaffine transformations. This is done by a least squares minimisation with respect to the $\mathbf{v}(x)$ (in this case computed using the LogDemons algorithm).

In [5], the authors showed that an incompressibility penalisation and a neighbouring region regularisation can be added to this model by penalising the trace of the gradient of the polyaffine velocity field, and by regularising the similarity between neighbouring regions via a distance term. Both terms were added in the least squares minimisation to obtain a penalised least squares error term, still linear in terms of M .

2.2 AHA Cardiac-Specific Region Definitions

Computing regional statistics at a population level requires a low-dimensional and consistent division of the ventricle. The American Heart Association (AHA) proposed a standardised method to segment the myocardium, which can be used to define 17 anatomically meaningful regions of the ventricle.

A binary mask can be semi-automatically created using a tool provided within the CardioViz3D package³. This mesh can then be divided into the 17 regions using a semi-automatic C++ segmentation tool. Since we are dealing with a small population, we can afford to do the segmentation semi-automatically. However, tools exist to do the segmentation automatically and could be used for large-scale implementation.

2.3 Polyaffine Weight Functions

The weights $\omega_i(x)$ can be defined by a simple Gaussian function as

$$\omega_i(x) = -\exp\left(\frac{\kappa}{2}(x - \bar{x}_i)^T \phi_i^{(-1)}(x - \bar{x}_i)\right), \quad (2)$$

with $\bar{x}_i$ the barycentre (centre point) of zone i and ϕ_i the corresponding covariance matrix as in [8]. This allows a smooth transition between neighbouring regions, which is necessary in order for the polyaffine fusion to remain diffeomorphic. The parameter κ controls how much influence neighbouring regions have on each other to allow independent transformations for $\kappa > 1$ towards smoother transitions between regions for $\kappa < 1$.

³ <http://www-sop.inria.fr/asclepios/software/CardioViz3D/>

3 Left-Ventricular Motion Tracking Validation on a Synthetic Sequence

In order to validate the accuracy of the registration, a data-set of 5 synthetic cine-MR sequences, with known ground truth were used from the data-set of synthetic image sequences generated by Prakosa et. al [9]. All sequences used for the validation were generated from a single patient sequence, by applying a bio-mechanical model to the reference image with different contraction parameters in order to simulate different motion dynamics.

3.1 Motion tracking validation and comparison to iLogDemons

Given the ground truth deformation from each frame to the reference frame (end diastolic frame), the L^2 -norm error between the computed deformation and the ground truth deformation was computed for each motion simulation over the sequence. The computed errors are shown in Fig. 1, comparing against the error of the iLogDemons algorithm [10] applied to the same sequences. From these curves we can see that the proposed method gives lower errors than the iLogDemons method, which may be because affine motion is more robust.




Fig. 1. Left: Total L^2 -norm error for the proposed method and the iLogDemons method per frame for the 5 simulated sequences. Right: The L^2 norm error for one case comparing the full polyaffine model (blue solid line), the model computed from only the first 3 modes (blue dashed line), and the iLogDemons algorithm (dashed red line).

3.2 Singular value decomposition on the polyaffine parameters

Given the low number of parameters needed to describe the deformation (204 per frame = 17 regions $\times$ 12 affine parameters), a singular value decomposition (SVD) of the parameters can be easily computed. For a single given case, the SVD of the affine matrix is computed as:

$$SVD(M) = U \cdot S \cdot V^T, \quad (3)$$

where U and V are unitary matrices, S is a diagonal matrix and V^T denotes the matrix transpose of V . The matrix M can be constructed by stacking the affine parameters for all the regions in a column for all the frames, so that the final matrix will be of size 204×29 .

Using three modes, more than 95% of the variability is captured. Based on this fact, a personalised reduced model was created using the first three modes of each subject. These were then used to re-create the deformation using the polyaffine model. The resulting error curve is shown in Fig. 1 (right) for a single case (similar curves can be shown for the other cases). As expected, there are slightly higher errors than the original polyaffine model, though the error is still lower than the iLogDemons algorithm. Therefore, we are able to describe the deformation by a reduced number of parameters - now 702 vs. 5916 for the full polyaffine model, and compared to voxel-based methods on the order of millions (image size in each of the x,y,z directions $\times 3 \times$ number of frames).

4 Group-wise Statistical Analysis of Motion

The motion tracking was applied to cine-MRI image sequences of 15 healthy volunteers obtained from the STACOM 2011 MICCAI cardiac motion tracking challenge database [1]. Steady-state free precision cine MRI were acquired using a 3T scanner (Philips Achieva System, Philips Healthcare) in the short axis view. Following the same protocol as in [1], the median error for the 15 volunteers was $3.55mm$, compared to $3.17mm$ for TDDFD and $4.78mm$ for iLogDemons. In order to be able to compare the parameters between subjects, rigid alignment to a common space was first performed using a simple landmark based registration.

4.1 Group-wise spatial mode decomposition

Given a set of subjects drawn from the same population (in this case healthy volunteers), we are interested in finding a group-wise spatial basis of the motion within this population. This can be done by using again SVD, but this time on the combined matrices of all the subjects (over all the regions, and all the time frames). By stacking the matrices column-wise, we can obtain a group-wise spatial basis, each with patient-specific time loadings (see Fig. 2). In this way, we can consider that the expected motion parameters for all the subjects are the same since they are drawn from the same population, but with different offsets (loadings) for the time.

Applying this to the 15 healthy subjects, we find that 3 modes were sufficient for describing 97% of the variability. Patient-specific models can be built using these first three group-specific spatial modes with the subject-specific temporal loadings.

4.2 Group-wise temporal mode decomposition

We may conversely be interested in finding a temporal basis for a given population. This could, for example, be useful in detecting abnormal motion by those with a temporal trigger significantly different from normal cases. Such a basis can be found by performing SVD on the group of affine matrices, this time stacked row-wise (see Fig. 3).




Fig. 2. Pipeline for computing a group-wise spatial basis by stacking the matrices for all the patients over all the regions and all time column-wise, and performing SVD on the big matrix.




Fig. 3. Pipeline for computing a temporal basis for a given population. The affine matrices of all the regions, parameters and frames for a given set of patients are stacked row-wise to perform a SVD.

By applying this decomposition to the 15 volunteers, we find that in this case 8 modes are needed to capture 90% of the variability. Given that the sequences were not fully temporally aligned, a larger number of modes is expected. Plotting the time signal for the first 8 modes only (Fig. 4, left), we find that the first mode corresponds well to the normal volume curve (Fig. 4, centre). Combining the first two modes has the effect of shifting the peak temporally (Fig. 4, right).

To test the generalisation of the temporal basis to a new subject, a leave-one-out study was performed on the 15 volunteers by learning the model on 14 volunteers and testing on the 15th. The residual error between the original affine matrices and the model computed from 3 modes was computed for the learning set to determine the range of expected residual values, then on the testing set to evaluate the generalisation of the basis (see Fig. 5 left). We observed that there is one outlier in the testing set, and went back to the data to analyse why this subject does not fit the normal model. Interestingly, we found that in fact this patient had trouble during the scan and moved significantly in the scanner due to discomfort. Therefore, the observed difference could in fact be due to motion artifacts. The model was then learned for the 14 normal volunteers (excluding the outlying volunteer) and projected onto two heart failure patients with known abnormal motion. The computed residual errors were much higher than the




Fig. 4. The first three temporal modes computed from the model reduction described in Fig. 3 (left) with estimated locations of the cardiac phases (black vertical lines). For comparison, the volume and electrocardiogram curves are shown (centre) (image modified from <http://en.wikipedia.org/wiki/File:WiggersDiagram.png>). The combination of the first two modes (right).

expected range (see Fig. 5 right), therefore the model is able to successfully distinguish abnormal motion in a quantitative manner.




Fig. 5. Left: Leave-one-out testing of the expected residual from using a 3-mode model rather than full polyaffine model with learning on 14 volunteers (left box-plot) and testing on the 15th volunteer (right box-plot). Right: Computed residual on the 3-mode model for 15 volunteers used in the learning (box-plot), and residual of the projection of the 3 modes on two pathological cases (blue stars).

5 Discussion and Future Work

The current work presents patient-specific, as well as group-wise analysis of motion by approximating the motion by a cardiac specific polyaffine model, then performing model reduction to the polyaffine parameters. Since the motion is described by a small number of meaningful parameters, the analysis can be done on a significantly reduced number of parameters and in a meaningful way, by decoupling the spatial and temporal aspects. By performing the model reduction in this way, rather than directly computing principal component analysis on the affine parameters at all time frames for all the subjects (vectorising the affine

and temporal components (represented by the green face of the cube of Fig. 3) for all subjects (represented by the orange face of the cube)), we de-couple the temporal and spatial components. We were able to detect irregular motion within the population, as well as distinguish abnormal dynamics from subjects outside the population using the temporal basis. For future work, temporally and spatially aligning the sequences first, as in [11], may give more meaningful bases with fewer modes needed to describe the variability. This work describes a proof of concept of the proposed method, which needs to be further validated on larger data-sets.

Acknowledgements: This project was partially funded by the Care4Me ITEA2 project and the Swiss National Science Foundation.

References

1. Tobon-Gomez, C., De-Craene, M., McLeod, K., Lautz, T., Shi, W., Hennemuth, A., Prakosa, A., Wang, H., Carr-White, G., Kapetanakis, S., Lutz, A., Rasche, V., Schaeffter, T., Butakoff, C., Friman, O., Mansi, T., Sermesant, M., Zhuang, X., Ourselin, S., Peitgen, H.O., Pennec, X., Razavi, R., Reuckert, D., Frangi, A., Rhode, K.: Benchmarking framework for myocardial tracking and deformation algorithms: An open access database. *Med. Image Anal. (MedIA)* (2013)
2. Chandrashekara, R., Rao, A., Sanchez-Ortiz, G., Mohiaddin, R., Rueckert, D.: Construction of a statistical model for cardiac motion analysis using nonrigid image registration. In: *Inf. Proc. Med. Imaging (IPMI)*, Springer (2003)
3. Duchateau, N., Craene, M.D., Piella, G., Silva, E., Doltra, A., Sitges, M., Bijmens, B.H., Frangi, A.F.: A spatiotemporal statistical atlas of motion for the quantification of abnormal myocardial tissue velocities. *Med. Image Anal. (MedIA)* **15**(3) (2011)
4. De Craene, M., Duchateau, N., Tobon-Gomez, C., Ghafaryasl, B., Piella, G., Rhode, K., Frange, A.: SPM to the heart: mapping of 4D continuous velocities for motion abnormality quantification. In: *Proc. ISBI*. (2012)
5. McLeod, K., Seiler, C., Sermesant, M., Pennec, X.: A near-incompressible poly-affine motion model for cardiac function analysis. In: *Proc. STACOM MICCAI Workshop. LNCS*, Springer (2012)
6. Metz, C.T., Baka, N., Kirisli, H., Schaap, M., Klein, S., Neefjes, L.A., Mollet, N.R., Lelieveldt, B., de Bruijne, M., Niessen, W.J., et al.: Regression-based cardiac motion prediction from single-phase CTA. *Trans. Med. Imaging (TMI)* **31**(6) (2012)
7. Hoogendoorn, C., Sukno, F.M., Ordás, S., Frangi, A.F.: Bilinear models for spatiotemporal point distribution analysis. *Int. J. Comp. Vis (IJCV)* **85**(3) (2009)
8. Seiler, C., Pennec, X., Reyes, M.: Capturing the multiscale anatomical shape variability with polyaffine transformation trees. *Med. Image Anal.* (2012)
9. Prakosa, A., Sermesant, M., Delingette, H., Marchesseau, S., Saloux, E., Allain, P., Villain, N., Ayache, N.: Generation of synthetic but visually realistic time series of cardiac images combining a biophysical model and clinical images. *IEEE Trans. Med. Imaging (TMI)* (2012)
10. Mansi, T., Pennec, X., Sermesant, M., Delingette, H., Ayache, N.: iLogDemons: A demons-based registration algorithm for tracking incompressible elastic biological tissues. *Int. J. Comp. Vision (IJCV)* **92**(1) (2011)
11. Perperidis, D., Mohiaddin, R.H., Rueckert, D.: Spatio-temporal free-form registration of cardiac MR image sequences. *Med. Image Anal. (MedIA)* **9**(5) (2005)