Learning A Global Descriptor of Cardiac Motion from a Large Cohort of 1000+ Normal Subjects

Wenjia Bai[1], Devis Peressutti[2], Ozan Oktay[1], Wenzhe Shi[1], Declan P O'Regan[3], Andrew P King[2], and Daniel Rueckert[1]

[1] Biomedical Image Analysis Group, Department of Computing, Imperial College London, UK
[2] Division of Imaging Sciences and Biomedical Engineering, King's College London, UK
[3] MRC Clinical Sciences Centre, Hammersmith Hospital, Imperial College London, UK

Abstract. Motion, together with shape, reflect important aspects of cardiac function. In this work, a new method is proposed for learning of a cardiac motion descriptor from a data-driven perspective. The resulting descriptor can characterise the global motion pattern of the left ventricle with a much lower dimension than the original motion data. It has demonstrated its predictive power on two exemplar classification tasks on a large cohort of 1093 normal subjects.

1 Introduction

Motion, together with shape, reflect important aspects of cardiac function. Detecting abnormal motion of the cardiac ventricles is of clinical interest for both diagnosis and prognosis [1]. To characterise the dynamics of cardiac motion, many clinical features have been proposed, such as the ejection fraction, segmental velocity, strain, strain-rate, time to peak velocity, time to peak strain, wall motion score index etc [1–3]. Apart from these, some statistical features have also been proposed, such as the cross correlation of intensity profiles across time [4] or the distance of the intensity distribution across time [5]. These indices are normally empirically defined and present an intuitive picture of the motion profile.

Apart from the empirical features, an alternative way is to describe the motion from a purely data-driven perspective and to learn a descriptor of the motion pattern from a large group of subjects. For example, in [6], a pixelwise Gaussian distribution of the velocity is learnt from a normal population and used for detection of motion abnormality. A challenge for cardiac motion analysis is that its data is normally of high dimension. Therefore, dimensionality reduction techniques are often used to discover the underlying structure of the data. For example, in [7–9], principal component analysis (PCA) and independent component analysis (ICA) have been proposed to learn the modes of myocardial shape variation at end-systole (ES). Whereas PCA assumes a Gaussian distribution

and performs linear dimensionality reduction, manifold learning is a non-linear technique which aims to preserve the local structure of the data. It has recently gained a lot of attention in the medical imaging community [10, 11].

In this work, we propose to learn a global descriptor of cardiac motion from segmental motion trajectories using dimensionality reduction techniques. The descriptor is not confined to a single voxel or a single segment, but instead it characterises the global motion pattern of the whole left ventricle (LV). We compare the performance of both PCA and Isomap manifold learning for dimensionality reduction. To demonstrate the value of the motion descriptor, we use it to predict the gender and age of a subject and evaluate the performance on a large data set of 1093 normal subjects.

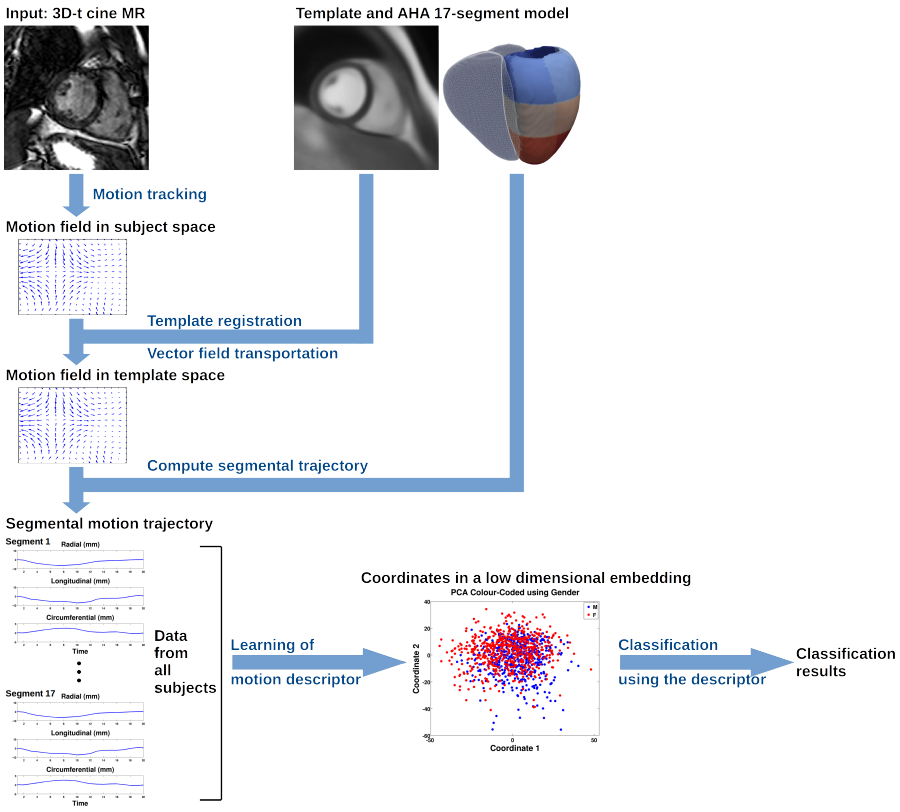

Fig. 1: The flowchart consists of motion tracking, spatial normalisation, descriptor learning and classification.

2 Methods

Prior to the learning of a motion descriptor, we first estimate motion from cine cardiac MR images. Since the heart of each subject lies at different locations and with different orientations, we perform spatial normalisation by registering and transporting all motion fields to a template space. The segmental motion trajectories are extracted and concatenated to form a high-dimensional feature vector. Dimensionality reduction is applied to the high-dimensional data leading to a global motion descriptor. Finally, we use the motion descriptor in exemplar classification tasks for gender classification and age prediction. Figure 1 illustrates the flowchart of the method and we will explain each step in the following.

2.1 Motion Tracking

In this work, we use cine MR for cardiac motion analysis. Other imaging modalities such as tagged MR or ultrasound (US) can also be used to capture the motion of the heart, which can provide different spatio-temporal resolution and image quality. The proposed motion descriptor is not confined to a specific imaging modality.

Motion tracking is performed for each subject using a 4D spatio-temporal B-spline image registration method with a sparseness regularisation term (TSFFD) [12]. The motion field estimate is represented by a displacement vector at each voxel and at each time frame t, which measures the displacement from the 0-th frame to the t-th frame. All the cine images in this work were acquired using the same imaging protocol, consisting of 20 time frames across a cardiac cycle with the 0-th frame representing the end-diastolic (ED) frame. Therefore, we do not perform temporal normalisation for the motion field.

2.2 Template Image and Spatial Normalisation

A template image is built by registering all the subject images at the ED frame and computing the average intensity image. In addition, the subject images are all segmented using a multi-atlas segmentation method [13]. The segmentation of the template image is then inferred by averaging all the subject segmentations. A template surface mesh is reconstructed from its segmentation and manually divided into 17 segments using the AHA model. The template and the segmental surface mesh are displayed at the top-right corner of Figure 1.

The motion field estimate lies within the space of each subject. To enable inter-subject comparison and analysis, all the subject images are aligned to the template image by non-rigid B-spline image registration [14]. Using the transformation between the template space and subject space, we transport the motion field of each subject to the template space. Let $x' = T(x)$ denote the transformation from the template space to the subject space, where x and x' are respectively the coordinates in the template space and in the subject space. By considering the spatial transformation as a change of coordinates, we have,

$$d(x,t) = J_{T^{-1}}(x')d'(x',t) \tag{1}$$

where d' denotes an infinitesimal displacement in the subject space, d denotes the corresponding infinitesimal displacement in the template space and $J_{T^{-1}}(x') \equiv \frac{dx}{dx'}$ denotes the Jacobian matrix of the inverse transformation.

2.3 Segmental Motion Trajectory

To characterise cardiac motion both spatially and temporally, we empirically define a high-dimensional feature vector using the segmental motion trajectory. S denotes the number of left ventricular segments. Since we use the AHA 17-segment model, $S = 17$. T denotes the number of time frames, which is equal to 20 for our data set. d denotes the dimension of the displacement vector and $d = 3$, which consists of radial, longitudinal and circumferential components. We compute the mean displacement for each segment at each time frame. The displacements across time for all the segments are concatenated to form the feature vector, which has the dimension of $S \times T \times d$ and contains information about the cardiac motion both spatially and temporally.

In principle, we can increase the spatial segments S so the feature vector describes more detailed motion at a higher spatial resolution. For example, we can compute the displacement for all the vertices of the myocardial mesh and concatenate them. However, we have found that it becomes computationally prohibitive to perform dimensionality reduction for vertex-wise motion data using techniques such as PCA. Also, since the cardiac motion is estimated using B-splines, displacements at neighbouring vertices are very similar and we may not need all the vertices to represent the motion data. Therefore, we adopt segment-wise motion data in this work.

2.4 Learning of a Motion Descriptor

Given the high-dimensional feature vector, we perform dimensionality reduction in order to find a descriptor which can characterise the motion with a low dimension. We compare two techniques, PCA and Isomap manifold learning [15]. The resulting low-dimensional coordinates are used as a motion descriptor.

PCA looks for a low-dimensional embedding of the data points that best preserves the variance. In the new coordinate system, the greatest variance of the data lies on the first coordinate, the second variance of the data on the second coordinate and so on. It is accomplished by eigen-decomposition of the data covariance matrix. In contrast to this, Isomap looks for a low-dimensional embedding that best preserves the geodesic distances between pairs of data points, i.e. the local data structure. It analyses the data structure as a graph, where each node denotes a data point and it is connected with K neighbours. The geodesic distances in the neighbourhood are preserved in the new coordinate system.

2.5 Application to Classification Tasks

To demonstrate the abundant information contained in the motion descriptor, we use the motion descriptor for two exemplar classification tasks, training SVM classifiers namely for gender classification and age prediction.

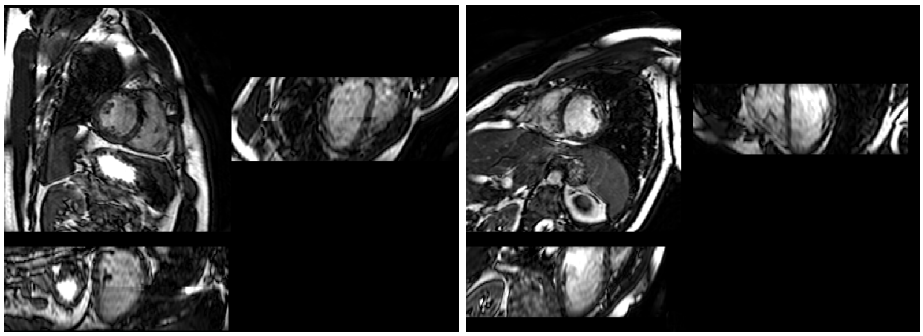

(a) Subject 1 (b) Subject 2

Fig. 2: Two exemplar cardiac MR images. Three orthogonal views are shown for each subject.

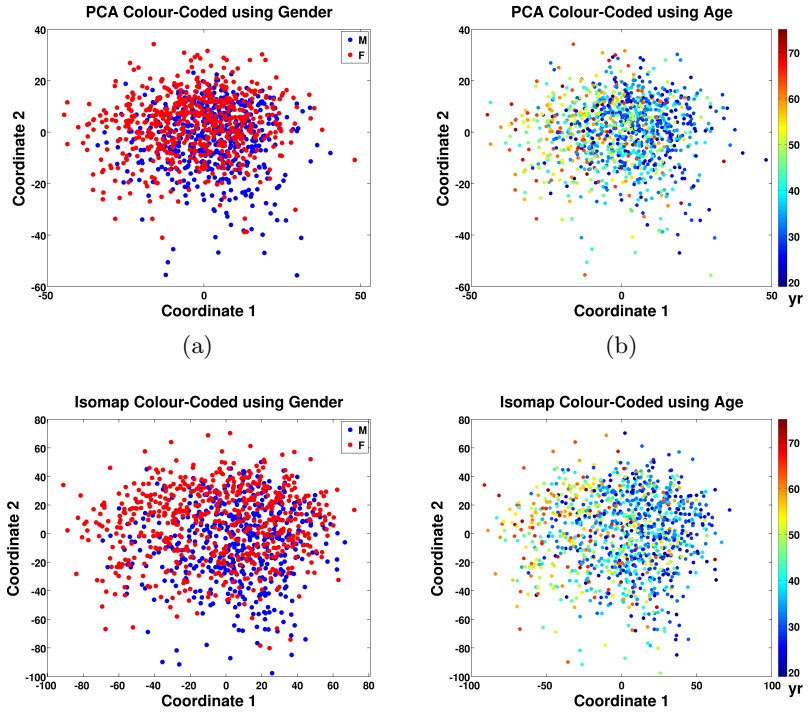

Fig. 3: Plot of the 1093 data points using the first two coordinates given by PCA (top row) or Isomap (bottom row). The data points are colour-coded using gender or age.

3 Experiments and Results

The data set used in this work consists of cardiac MR images of 1093 normal subjects (493 males, 600 females; age range 19-75 yr, mean 40.1 yr), which forms part of the UK 1000 Cardiac Phenomes project Cardiac MR was performed on a 1.5T Philips Achieva system (Best, Netherlands). The maximum gradient strength was 33 mT/m and the maximum slew rate 160 mT/m/ms. A 32 element cardiac phased-array coil was used for signal reception. Scout images were obtained and used to plan a single breath-hold 3D cine balanced steady-state free precession (b-SSFP) images in the left ventricular short axis (LVSA) plane from base to apex using the following parameters: repetition time msec/echo time msec, 3.0/1.5; flip angle, 50°; bandwidth, 1250 Hz/pixel; pixel size 2.0 × 2.0 mm; section thickness 2 mm overlapping; reconstructed voxel size, 1.25 × 1.25 × 2 mm; number of sections, 50 - 60; cardiac phases, 20; sensitivity encoding (SENSE) factor, 2.0 anterior-posterior and 2.0 right-left direction. Two exemplar images are displayed at Figure 2.

We extracted motion descriptors using PCA or Isomap from this data set. Figure 3 shows the first coordinates given by PCA and Isomap. The data points are colour-coded using gender or age. Figure 3(a) shows that the male subjects are more likely to be distributed at the bottom-right corner using the PCA coordinates, whereas the female subjects are more likely to be at the top-left corner. Figure 3(b) shows that the age of the subjects follows a right-to-left trend, gradually changing from young to old. The Isomap coordinates in Figure 3(c) and (d) reflect a similar trend as the PCA coordinates. This means that the motion descriptor, though in a low dimension, contains abundant information for motion data analysis.

SVM classifiers were trained using the segmental trajectories or the motion descriptors as input. We used the radial basis function (RBF) as the kernel and the default parameter settings in R. Ten-fold cross-validation was performed for performance evaluation. The performance was evaluated using the percentage of correct gender classification and the age prediction error. There is one parameter for Isomap, which is the number of neighbours K. We tuned the parameter and found that $K = 20$ achieved the best performance.

Table 1 lists the classification performance using the original segmental trajectories and using the motion descriptors with different dimensions. The original dimension of the feature is $S \times T \times d = 1020$. It shows that the best accuracy is achieved using the original high-dimensional feature vector, which is over 89% accuracy rate for gender classification and only -0.13 yr error for age prediction. However, with a much lower dimension of only 100, PCA can also achieve very high accuracy, with over 87% accuracy rate for gender classification and -0.36 yr error for age prediction. More dimensions in PCA do not necessarily improve the performance, since the first few coordinates of PCA have already encoded most of the data variance.

For Isomap manifold learning, it achieves worse gender accuracy rate of 80.15% and age prediction error of -0.47 yr, when the same dimension of 100 is used. This may hint that the high-dimensional feature vectors are located in a

Table 1: Classification performance using the original segmental trajectories and using the motion descriptors with different dimensions. Gender classification is evaluated using the percentage of correct classification and age prediction is evaluated using the prediction error.

	Dim	Gender	Age (yr)		Dim	Gender	Age (yr)
Original	1020	89.94%	-0.13±7.09				
PCA	10	76.21%	-0.65±9.27	Isomap	10	76.30%	-1.03±9.82
	50	84.63%	-0.50±8.12		50	80.60%	-0.60±9.36
	100	87.38%	-0.36±7.77		100	80.15%	-0.47±9.47
	150	86.74%	-0.30±8.03		150	78.78%	-0.47±9.69
	200	86.46%	-0.16±8.42		200	77.68%	-0.50±9.94

relatively flat manifold and therefore do not need non-linear techniques for dimensionality reduction. We have also tested two other manifold learning methods, namely locally linear embedding (LLE) and Laplacian eigenmaps. They achieve similar or slightly worse performance than Isomap.

4 Discussion and Conclusions

In this work, we learn a motion descriptor completely from a data-driven perspective on a large population of subjects by looking for a low-dimensional embedding which can explain either the data variance (PCA) or the local data structure (Isomap). We have demonstrated the resulting motion descriptor can be useful for both data visualisation and classification tasks.

There are mainly two reasons that motivate us to use the dimensionality reduction techniques. First, it allows convenient visualisation of high-dimensional data so that we can appreciate the data distribution in a better way. Second, it avoids the curse of dimensionality that may occur during data analysis. However, the SVM classifier seems to be well adapted to high-dimensional data. This explains why classification using the original high-dimensional data also yields a very good performance, as shown in Table 1.

In this work, segmental displacement trajectory is used as a representation of cardiac motion. However, other representations such as velocity, strain and diastolic filling rate can characterise motion in different ways. Our future work includes estimation of myocardial strain and potentially segmental strain can also be concatenated into the feature vector and used for learning a motion descriptor.

Registration is a key step to normalise the motion field of each subject into a template space. Two measures are taken to reduce the impact of potential registration errors on subsequent motion analysis. First, six landmarks are used to initialise the image registration so that severe registration error is less likely to happen. Second, the mean motion trajectory is computed for each segment, which is more robust than vertex-wise motion trajectory.

Although we only show exemplar applications for gender and age prediction on a data set of normal subjects, potentially we can also apply the motion descriptor to other tasks which require motion as input. In the future, we plan to explore its application on cardiac patient classification, for example, to classify between CRT respondents and non-respondents by combining the motion descriptor with other information such as QRS duration, myocardial scar amount etc, which may together reveal interesting findings about cardiac diseases.

References

1. V. Mor-Avi, R.M. Lang, L.P. Badano, M. Belohlavek, et al. Current and evolving echocardiographic techniques for the quantitative evaluation of cardiac mechanics. *European Journal of Echocardiography*, 12(3):167–205, 2011.
2. R.M. Lang, M. Bierig, R.B. Devereux, F.A. Flachskampf, et al. Recommendations for chamber quantification. *European Journal of Echocardiography*, 7(2):79–108, 2006.
3. J. Garcia-Barnés, D. Gil, L. Badiella, A. Hernandez-Sabate, et al. A normalized framework for the design of feature spaces assessing the left ventricular function. *IEEE Transactions on Medical Imaging*, 29(3):733–745, 2010.
4. Y. Lu, P. Radau, K. Connelly, A. Dick, and G. Wright. Pattern recognition of abnormal left ventricle wall motion in cardiac MR. In *MICCAI*, pages 750–758. Springer, 2009.
5. M. Afshin, I. Ben Ayed, K. Punithakumar, M. Law, et al. Regional assessment of cardiac left ventricular myocardial function via MRI statistical features. *IEEE Transactions on Medical Imaging*, 33(2):481–494, 2014.
6. N. Duchateau, M. De Craene, G. Piella, E. Silva, et al. A spatiotemporal statistical atlas of motion for the quantification of abnormal myocardial tissue velocities. *Medical Image Analysis*, 15(3):316–328, 2011.
7. E. Remme, A.A. Young, K.F. Augenstein, B. Cowan, and P.J. Hunter. Extraction and quantification of left ventricular deformation modes. *IEEE Transactions on Biomedical Engineering*, 51(11):1923–1931, 2004.
8. K.Y.E. Leung and J.G. Bosch. Local wall-motion classification in echocardiograms using shape models and orthomax rotations. In *FIMH*, pages 1–11. Springer, 2007.
9. A. Suinesiaputra, A.F. Frangi, T. Kaandorp, H.J. Lamb, et al. Automated detection of regional wall motion abnormalities based on a statistical model applied to multislice short-axis cardiac MR images. *IEEE Transactions on Medical Imaging*, 28(4):595–607, 2009.
10. K.K. Bhatia, A. Rao, A. Price, R. Wolz, J. Hajnal, and D. Rueckert. Hierarchical manifold learning for regional image analysis. *IEEE Transactions on Medical Imaging*, 33(2):444–461, 2014.
11. D.H. Ye, B. Desjardins, J. Hamm, H. Litt, and K. Pohl. Regional manifold learning for disease classification. *IEEE Transactions on Medical Imaging*, 33(6):1236–1247, 2014.
12. W. Shi, M. Jantsch, P. Aljabar, L. Pizarro, W. Bai, et al. Temporal sparse free-form deformations. *Medical Image Analysis*, 17(7):779–789, 2013.
13. W. Bai, W. Shi, D.P. O'Regan, T. Tong, et al. A probabilistic patch-based label fusion model for multi-atlas segmentation with registration refinement: application to cardiac MR images. *IEEE Transactions on Medical Imaging*, 32(7):1302–1315, 2013.

14. D. Rueckert, L.I. Sonoda, C. Hayes, D.L.G. Hill, M.O. Leach, and D.J. Hawkes. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.
15. J.B. Tenenbaum, V. De Silva, and J.C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.