



Learning Associations Between Clinical Information and Motion-Based Descriptors Using a Large Scale MR-derived Cardiac Motion Atlas

Esther Puyol-Antón¹(✉), Bram Ruijsink^{1,2}, Hélène Langet³, Mathieu De Craene³, Paolo Piro⁴, Julia A. Schnabel¹, and Andrew P. King¹

¹ School of Biomedical Engineering and Imaging Sciences,
King's College London, London, UK
esther.puyol_anton@kcl.ac.uk

² St Thomas' Hospital NHS Foundation Trust, London, UK

³ Philips Research, Medisys, Paris, France

⁴ Sanofi R&D, Paris, France

Abstract. The availability of large scale databases containing imaging and non-imaging data, such as the UK Biobank, represents an opportunity to improve our understanding of healthy and diseased bodily function. Cardiac motion atlases provide a space of reference in which the motion fields of a cohort of subjects can be directly compared. In this work, a cardiac motion atlas is built from cine MR data from the UK Biobank (≈ 6000 subjects). Two automated quality control strategies are proposed to reject subjects with insufficient image quality. Based on the atlas, three dimensionality reduction algorithms are evaluated to learn data-driven cardiac motion descriptors, and statistical methods used to study the association between these descriptors and non-imaging data. Results show a positive correlation between the atlas motion descriptors and body fat percentage, basal metabolic rate, hypertension, smoking status and alcohol intake frequency. The proposed method outperforms the ability to identify changes in cardiac function due to these known cardiovascular risk factors compared to ejection fraction, the most commonly used descriptor of cardiac function. In conclusion, this work represents a framework for further investigation of the factors influencing cardiac health.

Keywords: Cardiac motion atlas · Non-imaging data · Dimensionality reduction · Multivariate statistics

1 Introduction

Although much is known about the factors influencing healthy and diseased cardiac development, the recent availability of large scale databases such as the UK Biobank represents an excellent opportunity to extend this knowledge. The wide

range of clinical information contained in such databases can be used to learn new associations with indicators of cardiac function. In this paper we propose a framework for performing such an investigation.

Left-ventricular (LV) function has been traditionally assessed using global indicators such as ejection fraction and stroke volume. The main drawback of such indicators is that they do not provide localised information about ventricular function, and so their diagnostic power is limited. Motion parameters such as displacements allow for a more spatially and temporally localised assessment of LV function. Statistical cardiac motion atlases have been proposed as a way of making direct comparisons between motion parameters from multiple subjects [1, 10]. We use the UK Biobank database to build a cardiac motion atlas from cine MR images using around 6000 subjects with no known cardiovascular disease. To ensure that the estimated motion parameters are accurate, two automated quality control strategies are proposed to reject subjects with insufficient image quality. We use motion descriptors extracted from the atlas to identify relationships between cardiac function and other clinical information.

Previously, several papers have tried to analyse the influence of clinical information such as smoking, alcohol and body weight on cardiovascular disease [5–7]. However, these studies were based on a limited number of factors and used relatively crude indicators of cardiac health, such as categorical variables indicating disease state. In contrast, we exploit the large scale atlas to make use of more localised motion parameters and also examine a wider range of clinical information.

Contributions: Our scientific contributions in this paper are threefold. First, we present, for the first time, a large scale cardiac motion atlas built from the UK Biobank database. Second, we propose an automated quality control approach to ensure the robustness of the extracted motion descriptors. Third, we use the atlas to learn associations between descriptors of cardiac function and non-imaging clinical information. This automatic framework represents an ideal tool to further explore factors influencing cardiac health.

2 Materials

The UK Biobank data set contains multiple imaging and non-imaging information from more than half a million 40–69 year-olds. From this data set, only participants with MR imaging data were considered for analysis. Furthermore, participants with non-Caucasian ethnicity, known cardiovascular disease, respiratory disease, diabetes, hyperlipidaemia, haematological disease, renal disease, rheumatological disease, malignancy, symptoms of chest pain or dyspnoea were excluded.

MR imaging was carried out on a 1.5T scanner (Siemens Healthcare, Erlangen, Germany) [11]. Short-axis (SA) stacks covering the full heart, and two orthogonal long-axis (LA) planes (2-chamber (2Ch) and 4-chamber (4Ch) views) were available for each subject (TR/TE = 2.6/1.10ms, flip angle = 80°).

In-plane resolution of the SA stack and LA images was 1.8mm, with slice thickness of 8mm and 6mm for SA and LA respectively. 50 frames were acquired per cardiac cycle (temporal resolution $\approx 14\text{--}24\text{ ms/frame}$). In addition, the following non-imaging parameters were used for the statistical analysis: age, smoking status (smoker/non-smoker), body mass index (BMI), body fat percentage (BFP), basal metabolic rate (BMR), hypertension and alcohol intake frequency (occasionally/regularly). Demographics of the analysed cohort are presented in Table 1.

Table 1. Baseline characteristics for the healthy cohort from the UK Biobank data set. All continuous values are reported as mean (SD), while categorical data are reported as number (percentage).

	Total
Number of participants	6002
Male gender (n(%))	3046 (50.8%)
Age(years)	61 (8)
Body mass index (BMI) (kg/m^2)	26.4 (4.2)
Body fat percentage (BFP) (%)	28.2 (7.7)
Basal metabolic rate (BMR) (KJ)	6650.1 (1324.9)
Smokers (n(%))	1011 (38.2%)
Regular alcohol intake (n(%))	1407 (48.9%)
Hypertension (n(%))	2903 (48.4%)
Ejection fraction (EF)(%)	59.2 (6.3%)

3 Methods

In the following Sections we describe the estimation and analysis of the atlas-based motion descriptors. Details of the atlas formation procedure are reported in Sect. 3.1. Section 3.2 details the quality control (QC) strategies proposed. Section 3.3 describes the different dimensionality reduction algorithms used to characterise cardiac function. Section 3.4 presents the statistical methods used to identify relationships between the motion descriptors and non-motion clinical information.

3.1 Motion Atlas Formation

This section outlines the procedure used to build the motion atlas, which is based on the work published in [1] and [10].

LV Geometry Definition. A fully-convolutional network (FCN) with a 17 convolutional layer VGG-like architecture was used for automatic segmentation of the LV myocardium and blood-pool at end-diastole (ED) for SA and LA slices

[2, 14]. Each convolutional layer of the network is followed by batch normalisation and ReLU, except the last one, which is followed by the softmax function. In the case of the SA stack, each slice is segmented independently, i.e. in 2D. From the segmentations, a bounding box was generated and used to crop the image to only include the desired field of view, improving pipeline speed and reducing errors in motion tracking. The MR SA and LA segmentations were used to correct for breath-hold induced motion artefacts using the iterative registration algorithm proposed in [14]. The motion-corrected LA/SA segmentations were fused to form a 3D smooth myocardial mask. An open-source statistical shape model (SSM) was optimised to fit to this LV binary mask following an initial rigid registration using anatomical landmarks [1]. The use of landmarks ensured that the LV mesh was aligned to the same anatomical features for each patient, and the endocardial and epicardial surfaces of the mesh were smoothly fitted to the myocardial mask, providing surface meshes with point correspondence for each of the subjects and a 17 AHA regions segmentation.

Motion Tracking. A 3D GPU-based B-spline free-form deformation (FFD) was used [13] to estimate LV motion between consecutive frames. Subsequently, the inter-frame transformations were estimated and then composed using a $3D + t$ B-spline to estimate a full cycle $3D + t$ transformation.

Spatial Normalisation. Similar to [10], we transform each mesh to an unbiased atlas coordinate system using a combination of Procrustes alignment and Thin Plate Spline transformation. We denote the transformation for each subject n from its subject-specific coordinate system to the atlas by ϕ_n .

Medial Surface Generation. A medial surface mesh with regularly sampled vertices (≈ 1000) was generated from the SSM epicardial and endocardial meshes using ray-casting and homogeneous downsampling followed by cell subdivision.

Motion Reorientation. The displacements (\mathbf{u}_n) for each subject n were reoriented into the atlas coordinate system under a small deformation assumption using a push-forward action: $\mathbf{u}_n^{atlas} = \mathbf{J}(\phi_n(\mathbf{r}_n))\mathbf{u}_n$ [10], where $\mathbf{J}(\cdot)$ refers to the Jacobian matrix, and \mathbf{r}_n are the positions of the vertices in the ED frame.

Transform to Local Coordinate System. For each subject n the displacements in the atlas coordinate system \mathbf{u}_n^{atlas} were projected onto a local cylindrical coordinate system \mathbf{x}_n^{atlas} , providing radial, longitudinal, and circumferential information. The long axis of the LV ED atlas medial mesh was used as the longitudinal direction.

3.2 Quality Control

To ensure that the atlas was formed only from subjects with high quality cine MR images and robust motion descriptors, two QC methods were implemented to

automatically reject subjects with insufficient quality or incorrect segmentations and/or motion tracking. These two methods are:

QC1. SA images with insufficient image quality were excluded using two deep learning based approaches. First the 3D convolutional neural network (CNN) described in [8] was used to automatically identify images with motion-related artefacts such as mistriggering, arrhythmia and breathing artefacts. Second, subjects with incorrectly planned 4Ch LA images were automatically excluded using the CNN described in [9]. This pipeline identified images containing the left ventricular outflow tract, which is a common feature of poorly planned 4Ch images. A total of 806 subjects were rejected using the QC1 step.

QC2. Our second QC measure aimed to detect segmentation and/or motion tracking errors. LV volume curves over the cardiac cycle were computed based on the LV geometry and motion fields. Based on these curves, an automatic pipeline was used to identify curves with unrealistic properties. The LV curve was divided into diastole and systole based on the ED and end-systolic (ES) points, where ED corresponds to the point with maximum volume, and ES corresponds to the point with minimum volume. These two curves were treated independently. In the systole part of the curve, the downslope should be smooth and continuous and therefore the first derivative should only contain one peak. Cases with more than one (automatically identified using the first and the second derivative of the LV curve) peak were therefore excluded. In the diastole part the curves were divided into rapid inflow, diastasis and atrial systole phases. In the rapid inflow phase the upslope should be smooth and continuous and the first derivative should only contain one peak. Cases with more than one peak were therefore excluded. In the diastasis phase the curve should be flat and with no variation, so the first and second derivatives were used to exclude any cases with changes of volume higher than 10%. Finally, in the atrial systole phase the upslope should be smooth and continuous and the first derivative should only contain one peak. Cases with more than one peak were therefore excluded. The QC2 step excluded a total of 64 cases.

3.3 Dimensionality Reduction

In this work, we propose to learn descriptors of cardiac function from the motion atlas using dimensionality reduction techniques. More specifically, three techniques were compared: Principal Component Analysis (PCA) [3], Locally Linear Embedding (LLE) [12], and stacked denoising autoencoders (SDA) [15]. To form the inputs to these three algorithms, the local displacements per AHA region were averaged and these were concatenated into a column vector such that for subject n , $\mathbf{x}^{atlas} \in \mathbb{R}^L$, where $L = 3AT$, T is the number of cardiac phases and A is the number of AHA regions (17) in the atlas medial surface mesh. The column vectors for each subject were combined to produce a matrix $X = [\mathbf{x}_1^{atlas}, \dots, \mathbf{x}_N^{atlas}]$, where N is the number of subjects. The matrix X is

the input for all of the proposed dimensionality reduction techniques, and the output of all of them is a matrix $D \in \mathbb{R}^{(d \times N)}$, where d is the size of the low dimensional space.

3.4 Statistical Analysis

Multivariate statistics were used to find relationships between the low-dimensional motion descriptors and non-imaging clinical information. For categorical variables (i.e. smoking status, hypertension and alcohol intake) a multivariate Hotelling's T^2 test (assuming unequal variance) [4] was used to test for differences between groups, and the p -value was recorded. For continuous variables (age, BMI, BFP, BMR) a multivariate regression was used to compute the R statistic that measures the degree of relationship between the data. The R statistic and associated p -values were recorded. All the recorded p -values were adjusted using Bonferroni's correction for multiple tests.

Furthermore, we analysed the impact of using a cardiac motion atlas compared to standard clinical global measures such as ejection fraction (EF), the current most used clinical global measure of cardiac function. To this end, the EF for each subject was computed from the tracked meshes. For categorical variables, similar to the previous case, a paired t -test (assuming unequal variance) was used to test for differences between groups, and the p -value was recorded. For continuous variables a linear regression was used to compute the R statistic and the associated p -value. As before, all the recorded p -values were adjusted using Bonferroni's correction for multiple tests.

4 Experiments and Results

For all of the proposed dimensionality reduction algorithms the number of dimensions d retained was optimised. In the case of PCA, d was selected to retain 99% of the variance; in the case of LLE a grid search algorithm was used to vary d between 2 and 256 and the number of neighbours s was fixed to 10. Then, d and s were selected to minimise the reconstruction error ϵ . Finally, SDA was applied with 5 layers with size $[L, 2000, 1000, 500, d]$, a learning rate of 0.001, a corruption level of 50%, and number of epochs 500. The value of d was optimised using a grid search between 2 and 256, and d was selected to minimise the reconstruction error ϵ . For the three dimensionality reduction algorithms, the optimised d were: $d_{PCA} = 932$; $d_{LLE} = 32$ and $\epsilon_{LLE} = 0.37$ mm; $d_{SDA} = 50$ and $\epsilon_{SDE} = 0.15$ mm.

Table 2 shows the computed p -values with Bonferroni's correction and R statistics for the different clinical data for the different dimensionality reduction algorithms, and also using only EF as a motion descriptor. Bold numbers show p -values less than 0.05 (i.e. that showed statistical significance at 95% confidence). The results show that for all of the dimensionality reduction methods there was no statistical significance for age and BMI, while BFP, BMR, smoking status, hypertension and alcohol consumption are shown to be associated with

cardiac function. Furthermore, the R statistics for BFP and BMR for all of the non-linear dimensionality reduction algorithms are around 60%, which indicates a clear linear relationship with the data. Using only EF we can see that there is no statistical significance with any of the clinical information, and the R values are lower compared to the use of the cardiac motion atlas. Non-linear dimensionality reduction algorithms, i.e. LLE and SDA show lower p -values and higher R statistics compared to linear PCA for all of the clinical information analysed. For BFP, BMR, smoking status, hypertension and alcohol consumption, LLE shows a better linear fitting, with slightly higher R statistics compared to SDA. However, the difference between the two methods is not significant.

Table 2. Association of clinical information with estimated motion descriptors from the atlas represented by computed p -values and R statistics using multivariate statistics. The p -values and R statistics for EF were computed using paired t-tests and linear regression respectively. Bold numbers show p -values after Bonferroni’s correction less than 0.05, which are considered to be statistically significant at 95% confidence.

	R statistics (%)				p -value			
	PCA	LLE	SDA	EF	PCA	LLE	SDA	EF
Age	48.45	49.03	56.11	39.67	0.31	0.29	0.28	0.38
BMI	45.19	57.47	59.83	25.87	0.20	0.16	0.18	0.31
BFP	48.22	64.20	64.63	23.98	0.05	0.03	0.05	0.68
BMR	55.67	64.60	55.84	21.96	0.05	0.01	0.04	0.21
Smoking					0.05	0.02	0.04	0.56
Alcohol					0.05	0.03	0.02	0.62
Hypertension					0.03	0.01	0.03	0.15

5 Discussion

In this work, we have demonstrated the derivation of motion atlas-based descriptors of cardiac function and investigated their association with a range of clinical information. This is the first time that a cardiac motion atlas has been formed from a large number of subjects from the UK Biobank, and this has enabled us to produce a framework for investigating factors influencing cardiac health. We show that our method was able to identify factors known to impact cardiac health (directly or indirectly), such as hypertension, smoking, alcohol intake and body composition. Contrary to previous smaller studies, age was not related to large changes in cardiac function. We were unable to assess whether the apparent lack of association was influenced by a selection bias due to anticipated better survival of healthy versus relatively unhealthy individuals, the relatively small spread in ages (45–70) compared to other studies, the clustering of most participants in the middle age range (55–65), or was a consequence of the dimensionality reduction algorithms. We will investigate these possibilities in future work.

In addition, we demonstrated that the use of a cardiac motion atlas has added value over using only global indicators such as EF. The proposed method represents an important contribution to furthering our understanding of the influences on cardiac function. We used the selected categories as an illustration of the capabilities of the atlas method. However, the univariate analysis that we performed does not capture the complex interrelation of health factors, age and cardiac function. Future work will focus on a deeper analysis of the impact of these parameters and will also include further clinical characteristics. In this work, displacements were used as a representation of cardiac motion. However, other representations such as velocity and strain can characterise motion in different ways. Our future work also includes extension of the motion atlas to incorporate strain parameters. In conclusion, this work represents an important contribution to furthering our understanding of the influences on cardiac function and opens up the possibility of learning the relationships between motion descriptors and further non-imaging information.

Acknowledgements. This work is funded by the Kings College London & Imperial College London EPSRC Centre for Doctoral Training in Medical Imaging (EP/L015226/1) and supported by the Wellcome EPSRC Centre for Medical Engineering at Kings College London (WT 203148/Z/16/Z). This research has been conducted using the UK Biobank Resource under Application Number 17806.

References

1. Bai, W., Shi, W., et al.: A bi-ventricular cardiac atlas built from 1000+ high resolution MR images of healthy subjects and an analysis of shape and motion. *Med. Image Anal.* **26**(1), 133–145 (2015)
2. Bai, W., et al.: Semi-supervised learning for network-based cardiac MR image segmentation. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017*. LNCS, vol. 10434, pp. 253–260. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66185-8_29
3. Hotelling, H.: Analysis of a complex of statistical variables into principal components. *J. Educ. Psychol.* **24**(6), 417 (1933)
4. Hotelling, H.: The generalization of student's ratio. In: Kotz, S., Johnson, N.L. (eds.) *Breakthroughs in Statistics*. SSS, pp. 54–65. Springer, New York (1992). https://doi.org/10.1007/978-1-4612-0919-5_4
5. Jitnarin, N., Kosulwat, V., Rojroongwasinkul, N., et al.: The relationship between smoking, body weight, body mass index, and dietary intake among Thai adults: results of the National Thai Food Consumption Survey. *Asia Pac. J. Public Health* **26**(5), 481–493 (2014)
6. Mukamal, K., Chiuve, S., Rimm, E.: Alcohol consumption and risk for coronary heart disease in men with healthy lifestyles. *Arch. Intern. Med.* **166**(19), 2145–2150 (2006)
7. Nadruz, W., Claggett, B., Gonçalves, A., et al.: Smoking and cardiac structure and function in the elderly: clinical perspective: the ARIC study (atherosclerosis risk in communities). *Circ. Cardiovasc. Imaging* **9**(9), e004950 (2016)

8. Oksuz, I., et al.: Deep learning using K-space based data augmentation for automated cardiac MR motion artefact detection. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11070, pp. 250–258. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00928-1_29
9. Oksuz, I., Ruijsink, B., Puyol-Antón, E., et al.: Automatic left ventricular outflow tract classification for accurate cardiac MR planning, pp. 462–465 (2018)
10. Peressutti, D., Sinclair, M., et al.: A framework for combining a motion atlas with non-motion information to learn clinically useful biomarkers: application to cardiac resynchronisation therapy response prediction. *Med. Image Anal.* **35**, 669–684 (2017)
11. Petersen, S., Matthews, P., Francis, J., et al.: UK biobank’s cardiovascular magnetic resonance protocol. *J. Cardiovasc. Magn. Reson.* **18**(1), 8 (2016)
12. Roweis, S., Saul, L.: Nonlinear dimensionality reduction by locally linear embedding. *Science* **290**(5500), 2323–2326 (2000)
13. Rueckert, D., Sonoda, L., et al.: Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Trans. Med. Imag.* **18**(8), 712–721 (1999)
14. Sinclair, M., Bai, W., Puyol-Antón, E., Oktay, O., Rueckert, D., King, A.P.: Fully automated segmentation-based respiratory motion correction of multiplanar cardiac magnetic resonance images for large-scale datasets. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10434, pp. 332–340. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66185-8_38
15. Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., Manzagol, P.A.: Stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion. *J. Mach. Learn. Res.* **11**, 3371–3408 (2010)