Manifold modeling of the beating heart motion

Paul Stroe, Xianghua Xie, Adeline Paiement

Department of Computer Science, Swansea University, United Kingdom

Abstract. Modeling the heart motion has important applications for diagnosis and intervention. We present a new method for modeling the deformation of the myocardium in the cardiac cycle. Our approach is based on manifold learning to build a representation of shape variation across time. We experiment with various manifold types to identify the best manifold method, and with real patient data extracted from cine MRIs. We obtain a representation, common to all subjects, that can discriminate cardiac cycle phases and heart function types.

1 Introduction

The purpose of this study is to identify the main modes of variation of the myocardium's shape during the cardiac cycle. Possible uses of such a model of heart deformation include assessing heart function for diagnosis making, surgery planning, and heart motion correction or tracking. In the past, some works have tracked points at the surface of the heart to account for the heart motion due to combined heartbeat and breathing movements, e.g. [18]. These methods do not attempt to understand the heart's deformation, and therefore may lack of detail and robustness for other parts of the heart. Other methods have tried to describe these motions, such as [10] which built a respiration model using Principal Component Analysis (PCA). We propose a model of heart deformation due to heartbeat alone, that can be combined to such a respiration model to fully represent the complexity of heart motion.

Although heart deformation seems to have a high number of degrees of freedom, we investigate whether it lies on a lower dimensional manifold. This approach has been successfully used to encode the shape and the appearance of both the heart [11] and highly deformable objects such as the human body [3], with the aim to support tasks such as pose recovery, reconstruction, or tracking.

Heartbeat is a nonlinear process, thus the PCA modeling of [10] may not be sufficiently flexible to encode it, and a nonlinear manifold approach may be more suitable. Nonlinear manifolds have been used successfully on similar tasks in the past. Isomap combined with image distance measures representing deformations in MR images provided a representation of non-rigid chest cavity deformation due to breathing [17]. In [11], the authors demonstrated the ability of Diffusion Map to learn the right ventricle's shape variations in order to drive its segmentation. The closest work to our own is that of Gifani et al. [5], who used Locally Linear Embedding (LLE) on 2D echocardiography images to model the cardiac cycle of individual subjects and to automatically detect its phase. In this

work, we also apply nonlinear manifold learning to the modeling of myocardium shape deformation during the cardiac cycle. However, we aim to build a more general model that is applicable to a larger population group, and we use 3D shape information rather than 2D images. To the best of our knowledge, this is the first study to attempt this general modeling of heart shape's deformations. We test and compare several popular manifold types, in order to determine the best fit to this motion.

In the remainder of this article, we introduce our methodology in Section 2, where we also present the different manifold types we investigated. We report and discuss our comparative results in Section 3, and conclude in Section 4.

2 Proposed Approach

We propose to model deformations of the myocardium using manifold learning. As illustrated in Fig. 1, we use temporal sequences of 3D myocardium shape data, reconstructed from cardiac cine MRIs. After spatial and temporal normalization, the deformations of these 3D shapes are encoded in Difference Volumes (DV), then captured in a manifold. Next, we detail in turn these different steps.

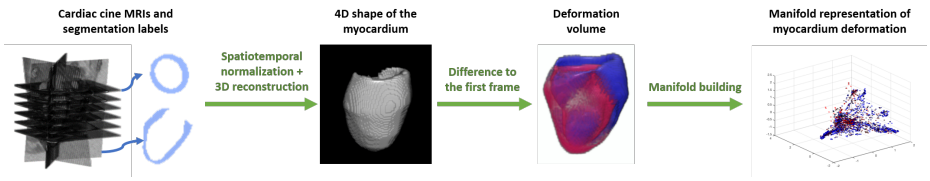

Fig. 1: Overview of proposed modeling of myocardium motion during heartbeat

2.1 Data

We illustrate our discussion using the CONSENSUS-AUTO dataset [7,4] that was released as part of the Cardiac Atlas Project (CAP). It provides cine cardiac MR images for 200 patients suffering from myocardial infarction and impaired LV contraction, with 19 to 30 time-frames per sequence. Semi-automated segmentation labels are also available for the myocardium in all slices and time-frames of 100 patients. We use these labels to visualize the left ventricle (LV) motion throughout the cardiac cycle. The MRIs are gated thus the respiratory motion does not affect the shape of the LV.

2.2 Preprocessing

In order to model the shape deformations of the myocardium, we start by recovering its 3D shape from the 2D segmentations of the MR images. This is done

using the IReSISD framework, introduced in [12] and [14] to jointly register, segment, interpolate 3D sparse MR volumes.

The sequences then need to be aligned temporally. Indeed, even though their acquisitions are ECG gated, we find that the volume curves of the LV blood pool are not perfectly synchronized (see Fig 2a). To correct this, we smooth the volume curve of each patient using moving average of 5 frames, then each sequence is translated rigidly to position their lowest volume frame at the start (Fig. 2b). This frame corresponds to the end-systole. Upon examination, we find that the maximum volume frames (end-diastole) are also well aligned by this method.

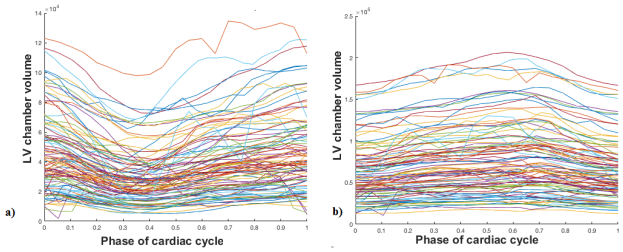

Fig. 2: Volume of the LV blood pool across the heart cycle. Note that initial curves (a) are not well aligned, while curves in (b) have been synchronized.

We then align the subjects spatially to compensate for different acquisition parameters and patient's position and morphology. First, we align the myocardium shapes at the end-diastole frame of each sequence using Procrustes alignment with translation and scaling. The rotation alignment along the LV's long axis is more difficult to perform using the LV shape only, because this shape tends to be quite symmetrical. Therefore, the second alignment step is based on the original MR images that contain both LV and RV. We use the middle SA slice of each patient's first frame, and select manually 14 landmarks around the right ventricle (RV) plus one in the middle of each ventricle (see Fig. 3). Overall, we found that the average, median, and maximum corrections to apply in our dataset are 16.12°, 13.54°, and 56.45°respectively.

2.3 Manifold Construction

We propose to model the deformations of the myocardium during the cardiac cycle using manifold learning. Heart shapes were found to form a non-linear manifold, e.g. in [11]. The beating motion of the heart is also complex and non-linear. Therefore it is reasonable to assume that heart deformations would form a manifold that is nonlinear. We experimented with a baseline linear PCA and several nonlinear manifold learning methods, and we will present their comparative results in Section 3.

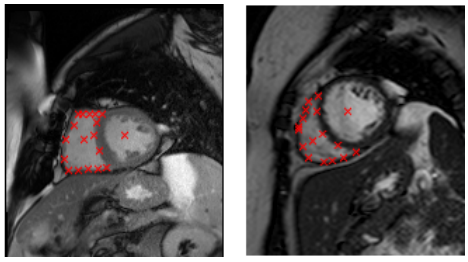

Fig. 3: Correction of rotation inconsistency between patients using landmarks (red crosses) alignment. The illustration here represents a rotation of 29°.

We build our manifold using deformation information encoded as the difference between a shape at time t and the shape of the same sequence at time t_0. We call this representation "Difference Volume" (DV), because it is built by taking the difference of the two volumes at t and t_0. In the DV, 0 values denote no change, values 1 are parts of the shape in t that appeared with regards to time t_0, and values -1 are parts that disappeared. An illustration is presented in Fig. 4. This representation has the advantage to preserve all the spatiotemporal information on deformation, while keeping the "shape-like representation" principle that was successful in previous works on shape manifold learning.

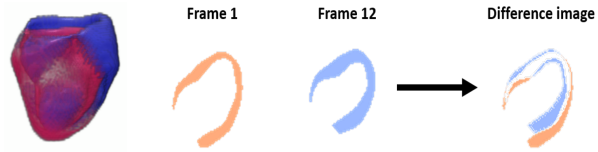

(a) (b)

Fig. 4: We encode the deformation information in a Difference Volume (DV) (a) that contains values 1 and -1 where the shape in frame t appears or disappears, respectively, with regards to initial shape at t_0. This is illustrated in 2D in (b).

We experimented with a number of manifold learning methods, which we describe briefly next. **Principal Component Analysis (PCA)** [6] finds the main uncorrelated orthogonal modes of variations in the data, to define a new space. **Diffusion Maps** [2] attempt to discover a lower-dimensionality space where distances reflect data similarity. Similarities are estimated as diffusion distances in the original space and computed by a random walk. We use the variant **Robust Diffusion Map (RDM)** presented in [13], that provides extra robustness to outliers. **Isomap** [16] uses a neighborhood graph in order to preserve the pairwise geodesic distances of the underlying manifold. **Laplacian eigenmap** [1] also uses a neighborhood graph that may be seen as a discrete approxima-

tion of the low-dimensional manifold that preserves distances. **t-Distributed Stochastic Neighbor Embedding (tSNE)** [8] also takes neighborhood into account to build a space where similar and dissimilar objects are modeled respectively by nearby and distant points. **Locally-Linear Embedding (LLE)** [15] considers a manifold as a serie of overlapping linear patches, under the assumption that, at a small enough scale, the manifold can be approximated as linear. This method therefore preserves the local properties of the data, and is good at finding nonconvex manifolds. We used the Isomap and LLE implementations of the Matlab Toolbox for Dimensionality Reduction [9].

3 Results

In our implementation, the DV introduced in Section 2 has 703,308 dimensions, which corresponds to the number of voxels in the 3D reconstruction of the myocardium. In this section we investigate which manifold type provides the most informative representation of LV motion during the cardiac cycle. We build models for myocardium and LV blood pool.

3.1 Comparison of manifold learning methods

We use all the data to build the manifolds in order to perform a visual exploration of the quality of dimensionality reduction offered by the different methods. We visualize in Figs. 5 and 6 the first three dimensions of all tested manifold learning methods, for the myocardium and LV chamber respectively. Color denotes the phase of the cardiac cycle, with blue for diastole, red for systole, and color saturation ranging from 0 at the beginning of the phase to 1 at the end.

The number of dimensions in these figures is chosen for convenience of visualization. The three first dimensions correspond to the main modes of variations discovered by the methods. Thus, they allow for an easy visual assessment of the discriminative power of the learned representations. The optimal number of dimensions for applications that call for a more detailed analysis shall be further explored in the future. For example, for classification of two different heart function types, we might find that three dimensions are enough (although that is still to be established). However, for a finer characterization of the heart function, more dimensions are likely to be needed. For our current purpose of initial visual examination, this drastically reduced dimensionality while still allowing to distinguish the evolution of the heart deformations in the most successful representations.

We observe that the PCA and RDM methods produce a better separation between the two systolic and diastolic cardiac phases (represented in red and blue respectively) and thus are more suitable for capturing the heart deformation information. Isomap, Laplacian Eigenmap, and LLE create manifolds made up of a central cluster and narrow branches. They do not separate well diastole from systole, and do not seem to provide a meaningful representation of heartbeat motion. tSNE produces a uniform sphere, which denotes a failure of the

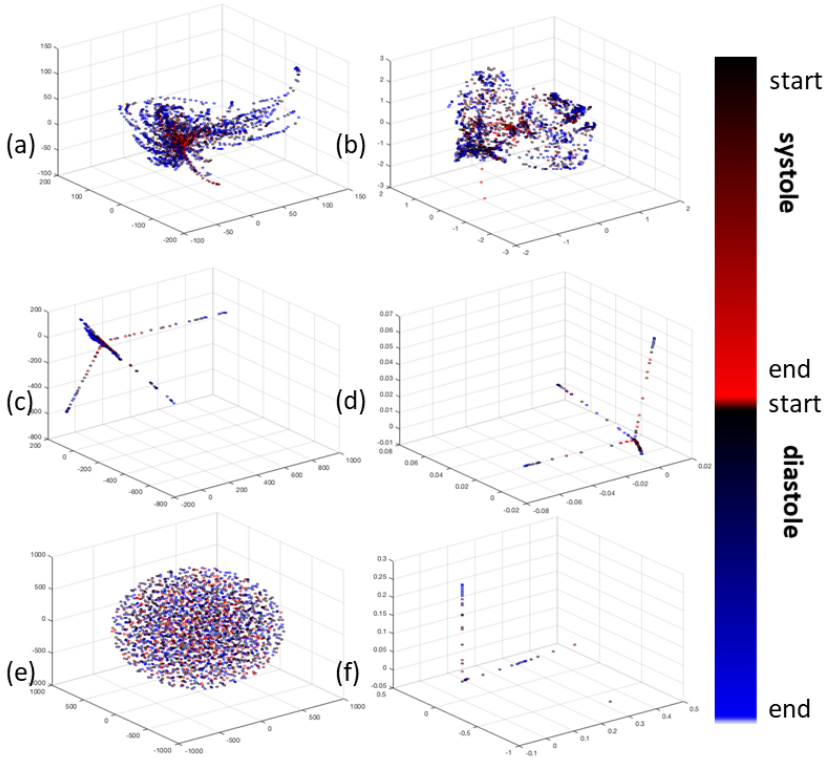

Fig. 5: Different manifold embedding of myocardium motion during cardiac cycle. The first three dimensions are visualized only. (a) PCA, (b) RDM, (c) Isomap, (d) Laplacian Eigenmap, (e) tSNE, (f) LLE. Legend: red and blue colors denote respectively systole and diastole, with saturation encoding progression of the phase, from beginning (black) to end (maximum saturation).

method to find a manifold. A possible explanation is that the optimum number of dimension for the tSNE representation is not three for this problem. In the rest of the article, we focus on the PCA and RDM manifolds that are identified as the most promising methods.

When looking at the PCA and RDM manifolds for myocardium and LV chamber, we notice that they all capture end-systole as a central cluster, while end-diastole has a larger area that denotes more variability in the heart deformations. This may be explained by the fact that we built the difference volumes by comparison against the end-systolic frame. Indeed, as can be seen in Fig. 2b, frame t_0 corresponds to the minimum of LV chamber volume.

The main difference between the two RDM manifolds, is that the end-diastole points seem more scattered in the myocardium manifold than in the LV chamber one. This suggests that the modes of deformation of the internal surface of the

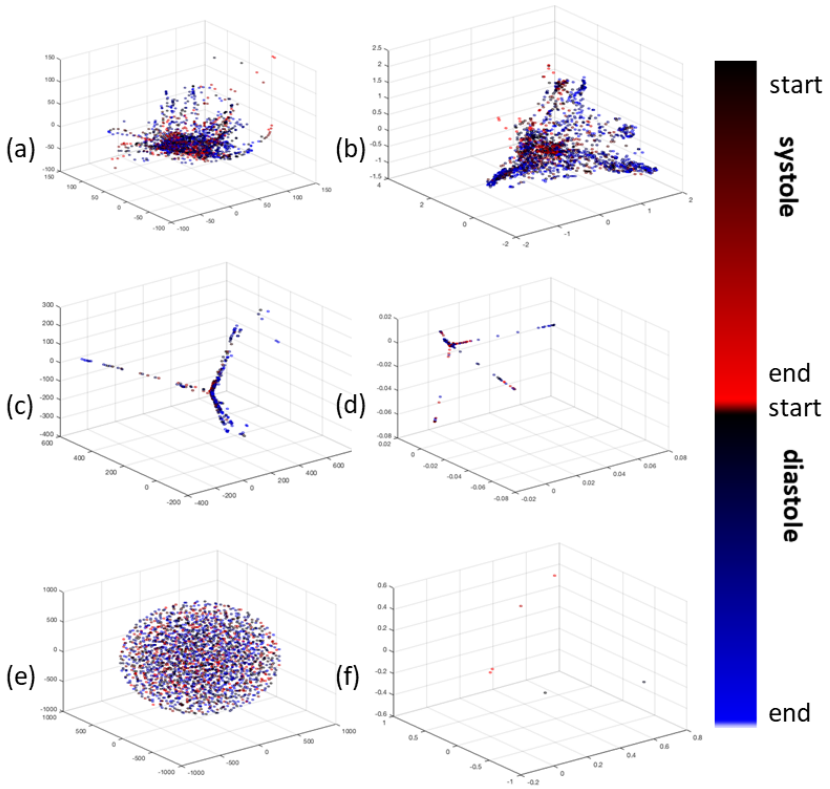

Fig. 6: Different manifold embedding of LV blood pool deformations during cardiac cycle. The first three dimensions are visualized only. (a) PCA, (b) RDM, (c) Isomap, (d) Laplacian Eigenmap, (e) tSNE, (f) LLE. Legend: red and blue colors denote respectively systole and diastole, with saturation encoding progression of the phase, from beginning (black) to end (maximum saturation).

myocardium, between the end-systolic and end-diastolic frames, is more consistent across subjects than the modes of deformation of the myocardium's external surface. We also note that the end-diastole points form two main branches in the LV chamber manifold, while two more localized clusters are visible in the myocardium manifold. This may indicate that the myocardium's internal surface presents a wider range of deformation intensities than its external surface. The two representations are otherwise of similar level of detail. We may use the LV chamber manifold when studying the blood pool deformations, and the myocardium manifold for a study of the myocardium itself.

PCA creates the second best manifold representation, which also identifies the systolic and diastolic phases. However, it embeds the systole shapes into a cloud of separate trajectories, in effect assigning regions of the manifold space

to individual patients. On the other hand, RDM groups similar shapes more successfully into two main motion modes rather than personal regions.

The two PCA and RDM methods will need to be further assessed in future works for given applications. Their optimal dimensionality would then need to be determined for that specific application.

3.2 Effect of data preprocessing on the manifold representation

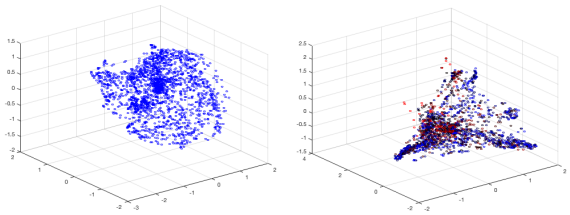

(a) (b)

Fig. 7: Effect of spatial and temporal normalization in the ability of the manifold representation to discriminate between phases of the cardiac cycle and cardiac function of the patients. (a) and (b) are with and without normalization. The color scheme is as in Figs. 5 and 6.

We compare the manifold created with data that underwent or not the data preparation steps presented in Section 2.2. As can be seen in Fig. 7a, the systolic and diastolic phases are less well separated. Also, in Fig. 7b, the manifold seems to identify two main modes of systole (which we attempt to interpret in Section 3.3), while in Fig. 7a the systole points form a fuzzy cloud. Our explanation for this effect is that the lack of normalization forces the manifold to also capture variations due to inconsistent spatial and temporal alignments. As a result, more dimensions would be needed to represent all relevant modes of variation.

3.3 Assessment of heart function types

We notice, in both RDM manifolds, two preferred modes of variation for the systole phase, denoted by the two branches or clusters of blue points. The dataset was made from patients that suffer from myocardial infarction. Thus, we hypothesize that these clusters may denote differences in scars within the myocardium. Unfortunately, the dataset does not identify patients groups, and this hypothesis is still to be verified with the help of a cardiologist's diagnosis. In the future, we plan to use data from healthy subjects in order to verify that anomalies are captured by our manifold in the same way as in [13].

4 Conclusion and Future Work

In this paper, we introduced a new methodology for modeling the temporal shape deformations of organs, and applied it to the myocardium motion during the cardiac cycle. We build our model from 4D shape data extracted from cine MRIs. The shape deformations were encoded in Difference Volumes then used to build a manifold representation.

We experimented with several popular manifold learning methods, on real patient data presenting mild to moderate anomalies affecting the beating heart's motion. We empirically determined that Robust Diffusion Map is best at encoding the range of shape deformations found across both the cardiac cycle and the patients' variable cardiac functions.

Our manifold model may be used in the future to characterize the motion, such as assessing heart function, including any anomaly type and intensity, for example by studying trajectories within the manifold space, as in [13]. It may also have uses in estimating the cardiac cycle phase from 3D shapes and possibly image volumes. When combined with a model of respiratory movements, this may be used for motion compensation. These applications will be the subjects of future works.

References

1. Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation. Neural Computation 15, 1373–1396 (2003)
2. Coifman, R.R., et al.: Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. Proc. National Academy of Sciences of USA 102(21), 7426–7431 (2005)
3. Elgammal, A., Lee, C.S.: The role of manifold learning in human motion analysis. In: Human Motion, pp. 25–56. Springer (2008)
4. Fonseca, C.G., et al.: The Cardiac Atlas Project An imaging database for computational modeling and statistical atlases of the heart. Bioinformatics 27, 2288–2295 (2011)
5. Gifani, P., et al.: Automatic detection of end-diastole and end-systole from echocardiography images using manifold learning. Physiological Measurement 31(9), 1091 (2010)
6. Jolliffe, I.T.: Principal Component Analysis. Springer, New York (2002)
7. Kadish, A.H., et al.: Rationale and Design for the Defibrillators to Reduce Risk by Magnetic Resonance Imaging Evaluation (DETERMINE) Trial. J Cardiovasc Electrophysiol 20, 982–987 (2009)
8. Maaten, L.v.d., Hinton, G.: Visualizing data using t-SNE. Journal of Machine Learning Research 9(Nov), 2579–2605 (2008)
9. Van der Maaten, L.: Matlab toolbox for dimensionality reduction
10. McClelland, J.R., et al.: Respiratory motion models: a review. Medical Image Analysis 17(1), 19–42 (2013)
11. Moolan-Feroze, O., et al.: Segmentation of the right ventricle using diffusion maps and markov random fields. In: MICCAI. pp. 682–689. Springer (2014)
12. Paiement, A., et al.: Integrated segmentation and interpolation of sparse data. IEEE TIP 23(1), 110–125 (2014)

13. Paiement, A., et al.: Online quality assessment of human movement from skeleton data. In: BMVC (2014)
14. Paiement, A., et al.: Registration and modeling from spaced and misaligned image volumes. IEEE TIP 25(9), 4379–4393 (2016)
15. Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. Science 290(5500), 2323–2326 (2000)
16. Silva, V.D., Tenenbaum, J.B.: Global versus local methods in nonlinear dimensionality reduction. In: Advances in neural information processing systems. pp. 705–712 (2002)
17. Souvenir, R., Pless, R.: Isomap and nonparametric models of image deformation. In: WACS-MOTIONS. vol. 2, pp. 195–200. IEEE (2005)
18. Yang, B., et al.: A triangular radial cubic spline deformation model for efficient 3D beating heart tracking. Signal, Image and Video Processing pp. 1–8 (2017)