



HAL
open science

Characterizing interactions between cardiac shape and deformation by non-linear manifold learning

Maxime Di Folco, Moceri Pamela, Patrick Clarysse, Nicolas Duchateau

► **To cite this version:**

Maxime Di Folco, Moceri Pamela, Patrick Clarysse, Nicolas Duchateau. Characterizing interactions between cardiac shape and deformation by non-linear manifold learning. *Medical Image Analysis*, 2021, pp.102278. 10.1016/j.media.2021.102278 . hal-03407714

HAL Id: hal-03407714

<https://hal.science/hal-03407714>

Submitted on 28 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Characterizing interactions between cardiac shape and deformation by non-linear manifold learning

Di Folco Maxime^{a,*}, Mocerì Pamela^b, Clarysse Patrick^a, Duchateau Nicolas^a

^aUniv Lyon, UCBL, Inserm, INSA Lyon, CNRS, CREATIS, UMR5220, U1294,69621 Villeurbanne, France

^bCentre Hospitalier Universitaire de Nice, Service de Cardiologie, Nice, France

Abstract

In clinical routine, high-dimensional descriptors of the cardiac function such as shape and deformation are reduced to scalars (e.g. volumes or ejection fraction), which limit the characterization of complex diseases. Besides, these descriptors undergo interactions depending on disease, which may bias their computational analysis. In this paper, we aim at characterizing such interactions by unsupervised manifold learning. We propose to use a sparsified version of Multiple Manifold Learning to align the latent spaces encoding each descriptor and weighting the strength of the alignment depending on each pair of samples. While this framework was up to now only applied to link different datasets from the same manifold, we demonstrate its relevance to characterize the interactions between different but partially related descriptors of the cardiac function (shape and deformation).

We benchmark our approach against linear and non-linear embedding strategies, among which the fusion of manifolds by Multiple Kernel Learning, the independent embedding of each descriptor by Diffusion Maps, and a strict alignment based on pairwise correspondences. We first evaluated the methods on a synthetic dataset from a 0D cardiac model where the interactions between descriptors are fully controlled. Then, we transferred them to a population of right ventricular meshes from 310 subjects (100 healthy and 210 patients with right ventricular disease) obtained from 3D echocardiography, where the link between shape and deformation is key for disease understanding. Our experiments underline the relevance of jointly considering shape and deformation descriptors, and that manifold alignment is preferable over fusion for our application. They also confirm at a finer scale the characteristic traits of the right ventricular diseases in our population.

Keywords: Cardiac imaging, dimensionality reduction, information fusion, manifold learning, myocardial strain

1. Introduction

The computational analysis of medical images allows going beyond the measurements performed in clinical routine, by considering high-dimensional descriptors that represent complementary aspects of diseases. Finding a data representation suited for clinical understanding across populations is highly challenging, and even more

when several descriptors are considered. In this paper, we aim at relating shape and deformation observations, which are key to characterize the cardiac function. For this purpose, we demonstrate the relevance of recent representation learning schemes not only to combine several high-dimensional descriptors, but also to characterize their interactions depending on individuals and disease.

1.1. Low-dimensional embedding of multiple descriptors by manifold learning

Manifold learning methods provide a low-dimensional representation of the data that also preserves the global or

*Corresponding author:
Email address: difolco@creatis.insa-lyon.fr (Di Folco Maxime)

local relations between samples, and therefore the structure of the data space. These dimensions represent the main directions encoding specific properties in the data (principal geodesics (Tenenbaum et al., 2000), modes of diffusion (Coifman et al., 2005), etc.) depending on the specific manifold learning method used. These techniques are therefore recommended to analyze the distribution of populations or subgroups, and the relations between samples according to well-controlled distances. Nonetheless, they suffer from the absence of explicit mappings to navigate between the high- and the low-dimensional spaces. In this sense, they differ from the recent variational autoencoders (VAE) (Kingma and Welling, 2013), which constrain the latent space to fit a specific distribution but do not explicitly consider local relations between samples. The variant β -VAE reduces this limitation by introducing a weighting factor to control the relative contribution of sample interactions and distribution fit (Higgins et al., 2016). In contrast, manifold learning explicitly optimizes the pairwise interactions between samples. This is an asset for interpreting latent spaces as in our application. Besides, the latent dimensions in VAE are not necessarily ordered by importance, contrary to more classical dimensionality reduction methods, while this could also facilitate the interpretations by clinicians. These limitations are counterbalanced by intrinsic encoding and decoding operations with the VAE framework.

From simple linear methods such as principal component analysis to non-linear methods that can apprehend more complex structures, manifold learning (Yan et al., 2007) has been used on many medical imaging applications in the last decade. Nevertheless, most of them only take into account a single descriptor at once, while complex diseases require considering different although partially correlated descriptors.

Manifold learning methods that can exploit several high-dimensional features and their potential interactions can be classified into two main categories: fusion and alignment strategies (Li et al., 2018).

Fusion methods look for a single latent representation from several input descriptors. For example, Multiple Kernel Learning (MKL) builds upon the unified framework of manifold learning, where an affinity matrix -which quantifies the similarities between pairs of samples- is estimated for each descriptor. MKL performs a linear combination of the affinity matrices from several

descriptors and optimizes both the mapping to the low-dimensional representation and the weights given to each descriptor (Lin et al., 2011). Similarity Network Fusion (Wang et al., 2014) uses a cross-diffusion process to iteratively merge two affinity matrices and reach a consensus between their corresponding descriptors. Extensions of the Gaussian process latent variable models also allow the fusion of manifolds from a Bayesian perspective, which is interesting to express dependencies between variables (Lawrence and Moore, 2007). However, the relative contribution of each descriptor is not explicitly considered.

In contrast, alignment methods aim to capture the relationship between the descriptors by aligning the different representations learnt for each descriptor, which therefore become comparable (Ham et al., 2005). Partial Least Squares (PLS) and Canonical Correlation Analysis (CCA) are two linear methods that maximize the covariance or the correlation between the descriptors in the low-dimensional latent space. However in our application, non-linear methods may be more suited to estimate the manifolds associated to cardiac shape and deformation, and prevent analysis errors (Duchateau et al., 2012). Besides, extending PLS and CCA to more than two descriptors may not be straightforward. In this sense, the framework of Multiple Manifold Learning (MML) is relevant (Clough et al., 2019; Lindenbaum et al., 2020), as a generalization of the strict alignment proposed in (Ham et al., 2005). It consists in building a larger affinity matrix whose blocks either represent the affinity between samples according to one descriptor (the diagonal blocks) or across several descriptors (the extra-diagonal blocks). The low-dimensional representation is therefore estimated according to each descriptor and their correspondences. This framework has been illustrated for the alignment of different datasets (Valencia-Aguirre et al., 2011; Lee et al., 2016), of different magnetic resonance slices (Clough et al., 2019; Baumgartner et al., 2017) or to find regional associations between anatomical structures (Benkarim et al., 2020). These works slightly differ regarding the definition of the affinity matrix, and in particular its extra-diagonal blocks, formulated through permutation operators (Lee et al., 2016), correlations (Valencia-Aguirre et al., 2011; Benkarim et al., 2020), or the graph structure of the datasets (Baumgartner et al., 2015; Clough et al., 2019).

Nevertheless, these works only aimed at aligning differ-

ent datasets that originate from the same manifold, such as 2D MRI slices for 3D stacking purposes (Baumgartner et al., 2017; Clough et al., 2019) or images of objects represented from similar sets of viewpoints (Valencia-Aguirre et al., 2011)). Our problem is slightly different, as we consider descriptors belonging to different manifolds whose (unknown) link exists but can vary depending on the samples. We propose to go further and demonstrate that a manifold alignment strategy is also relevant to characterize the distribution of each descriptor in a population, while controlling the strength of the alignment for better disease-specific and subject-specific analysis.

1.2. Analyzing myocardial shape and deformation

In this work, we aim at exploring the link between two relevant descriptors of the cardiac function (myocardial shape and deformation) on complex diseases of the right ventricle, for which the overall ventricle is dilated but shape and deformation are affected differently depending on the disease, its severity, and subject-specific characteristics. Although myocardial shape and deformation correspond to high-dimensional information accessible at each point of the ventricle and each instant of the cycle, clinical practice and even research generally reduce it to scalars such as volumes, ejection fraction or global strains (Cikes and Solomon, 2016). This represents a strong truncation of the available data which could hamper disease understanding and patient management.

The low-dimensional embedding of myocardial shape or motion/deformation has been thoroughly explored in research, as reviewed in (Gilbert et al., 2020; Duchateau et al., 2020). Nonetheless, shape and deformation should not be considered independently from each other. They may undergo structural (Bijnens et al., 2012) and disease-related interactions as illustrated for our application of interest: the characterization of right ventricular diseases with pressure- or volume-overload (Moceri et al., 2018; Sanz et al., 2019). Details on the studied diseases and their specific shape-deformation abnormalities are given in Sec. 2.1

Few works investigated myocardial shape or deformation in the context of multiple high-dimensional imaging descriptors. For instance, (Sanchez-Martinez et al., 2017) combined multiple features among which myocardial velocities using non-linear techniques (MKL) to characterize different patterns of functional responses to stress

in the heart failure with preserved ejection fraction syndrome. The different descriptors may be partially correlated, but this aspect was not explicitly considered. The correlation between two descriptors (motion obtained from two different modalities) was specifically addressed in (Puyol-Antón et al., 2017) using linear embedding techniques such as PLS and CCA. In this case, the correlation between descriptors is high, as they correspond to two instances of the same objects seen with different modalities. In preliminary work, we tested these linear techniques on myocardial shape and deformation, but only limited our assessment to the observation of substantial differences in zones of high interest for the studied diseases (Di Folco et al., 2020). In the present paper, we investigate this link much more in detail, on both synthetic and (larger) real populations. Besides, we focus on non-linear methods, which are better suited to the underlying structure of the spaces to which shape and deformation descriptors belong, and provide better flexibility to study their interactions.

1.3. Proposed approach and contributions

We specifically address two challenges in this paper:

1. Finding a joint embedding for two high-dimensional and heterogeneous descriptors (myocardial shape and deformation) whose link depends on individuals and disease.
2. Characterizing this link on a relatively large database of Right Ventricular (RV) meshes from 3D echocardiography, from patients with pressure or volume overload, which may exhibit a dilated RV but for which local shape and deformation patterns, and their potential interactions, are key to assess the disease severity and evolution.

To do so, we propose to apply unsupervised manifold learning and especially MML to apprehend this link. We evaluated the interest of this alignment strategy over fusion methods (MKL), and the embedding of each descriptor independently (Diffusion Maps (DM)). We also evaluated its relevance over linear alignment based on PLS, which we used previously in (Di Folco et al., 2020).

For understanding purposes, we first tested our approach on a synthetic dataset of pressure and volume

curves from a 0D circulation model, where the link between two descriptors is fully controlled. Then, we evaluated the potential of this approach to characterize the link between cardiac shape and cardiac deformation on a population of patients with RV volume- or pressure-overload, from two 3D echocardiographic studies (Moceri et al., 2018, 2020), which exhibit both shape and deformation changes with disease. To do so, we proposed a set of experiments to explore the consistency between neighboring samples, the latent dimensions and the main anatomical characteristics encoded. Finally, we illustrated the interest of our approach by linking the results of the experiments to clinical observations.

2. Methods

2.1. Data

We processed a database of 310 subjects, whose RV endocardial surface was tracked with 3D echocardiographic image sequences using the commercial software 4D-RV Function 2.0 (TomTec Imaging Systems GmbH, Unterschleißheim, Germany). The meshes (valves excluded) consisted of 822 points and 1587 triangular cells. Point-to-point correspondences between the meshes from different subjects are ensured by the software, which allows straightforward inter-subject comparisons at each point or cell of the RV mesh. The database is composed of four different groups detailed in Table 1. The 100 cases in the control group served as reference for the shape meshes and for normality to quantify the anomalies of the diseased groups. In particular, 55 controls are age- and sex-matched with the Atrial Septal Defect (ASD) and Tetralogy of Fallot (ToF) groups.

ASD and ToF are two congenital diseases. Patients with ASD exhibit a hole in the inter-atrial septum. RV volume overload is observed when the hole becomes hemodynamically significant, because of chronic left-to-right shunt with increased pulmonary blood flow. Patients with ToF typically present a deviated conal septum, an overriding aorta with ventricular septal defect and pulmonary stenosis (muscular and valvar pulmonary stenosis). Repair usually occurs during infancy and the surgeon has to close the ventricular septal defect and relieve the RV outflow obstruction, thus being classically responsible for the development of pulmonary regurgitation that can impact

Table 1: Clinical characteristics of the studied subjects

Subgroup	Overload	RVEDV ¹ (mL)	RVEF ² (%)	Global RVAS ³ (%)	N
Controls	None	64.0 ± 25.3	54.2 ± 6.3	-32.5 ± 5.2	100
PH	Pressure	103.1 ± 49.2	34.9 ± 11.1	-20.8 ± 6.8	155
ASD	Volume	153.6 ± 91.2	55.2 ± 8.7	-35.4 ± 7.3	27
ToF	Volume	152.1 ± 61.8	43.6 ± 9.5	-26.2 ± 5.5	28
Total					310
Controls vs. PH ⁴		<0.001	<0.001	<0.001	
Controls vs. ASD vs. ToF ⁵		<0.001	<0.001	<0.001	

¹ Right Ventricle End-Diastolic Volume (RVEDV)

² Right Ventricle Ejection Fraction (RVEF)

³ Right Ventricle Area Strain (RVAS)

⁴ p-value from unpaired t-test

⁵ p-value from one-way ANOVA

the long term prognosis. The chronic pulmonary regurgitation induces RV volume overload and can be responsible of symptoms and lead to chronic heart failure. Both groups present RV volume overload with shape and deformation abnormalities, at different grades (Moceri et al., 2020) and different stages of the disease (Sanz et al., 2019). In contrast, patients with Pulmonary Hypertension (PH) undergo right ventricular pressure overload, which affects the RV contractility to preserve the cardiac output and leads to shape and deformation abnormalities (Moceri et al., 2018). For all, the clinical challenge is to be able to characterize the differences between these types of overload up to local abnormalities in the shape and deformation patterns.

2.2. Shape and deformation features

To reduce the bias due to different spatial heart positions, we realigned the meshes using generalized Procrustes analysis (Gower, 1975) with a similarity transform. For this population, we represent myocardial shape and deformation by two high-dimensional descriptors available at each point of the RV mesh. The first one, called area strain (Fig. 1a), stands for the relative change in the area of each mesh cell (in %) between End-Diastole (ED) and End-Systole (ES). In our implementation, area strain was redefined at each point of the mesh by averaging its values from the immediate adjacent cells. This provides slight smoothing of the strain patterns and therefore

reduces local artifacts due to the meshing, and enables analyzing deformation and shape at the same locations. Area strain therefore corresponds to a high-dimensional descriptor of dimensions 822×1 (822 is the number of points for each mesh). The second descriptor characterizes shape at ED, through a 3D displacement vector (Fig. 1b) computed at each mesh point between a given case and the average shape of the controls (obtained by Procrustes analysis on this subgroup), used as reference for normality. This descriptor is of dimensions 822×3 .




Figure 1: (a) Area strain (%) at ES for a healthy case. (b) Distance vector (mm) colored by magnitude for the same healthy case. The red mesh stands for the average shape of the controls, used as reference for the computations.

2.3. Manifold learning

Given a population of K subjects, let's denote $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_K]^T \in \mathbb{R}^{K \times M}$ (the dimensionality M is usually large) the input data for one descriptor. Manifold learning estimates a set of low-dimensional coordinates $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_K]^T \in \mathbb{R}^{K \times N}$ such that $N \ll M$.

2.3.1. One descriptor - Diffusion Maps

The unified framework for manifold learning (Yan et al., 2007) performs spectral embedding as in Laplacian eigenmaps. It constructs an affinity matrix $\mathbf{W} = [W_{ij}] \in \mathbb{R}^{K \times K}$ where:

$$W_{ij} = \exp \frac{-\|\mathbf{x}_i - \mathbf{x}_j\|^2}{\sigma^2} \quad (1)$$

where σ is the width of the kernel. Then, it minimizes:

$$\Phi(\mathbf{Y}) = \frac{1}{2} \sum_{i,j=1}^K \|\mathbf{y}_i - \mathbf{y}_j\|^2 W_{ij} = \text{tr}(\mathbf{Y}^T \mathbf{L} \mathbf{Y}) \quad (2)$$

s.t. $\mathbf{Y}^T \mathbf{D} \mathbf{Y} = \mathbf{I}$

where $\mathbf{D} = [D_{ij}]$ is a diagonal matrix such that $D_{ii} = \sum_j W_{ij}$. \mathbf{L} stands for the graph Laplacian and is defined by: $\mathbf{L} = \mathbf{D} - \mathbf{W}$. Intuitively, the cost function in Eq. 2 means that close samples \mathbf{x}_i and \mathbf{x}_j in the high-dimensional space (W_{ij} close to 1) will result in close low-dimensional coordinates \mathbf{y}_i and \mathbf{y}_j . Conversely, distant samples (W_{ij} close to 0) will not contribute to bring \mathbf{y}_i and \mathbf{y}_j closer.

Equation 2 amounts to solving the generalized eigenvalue problem $\mathbf{L} \mathbf{f} = \lambda \mathbf{D} \mathbf{f}$, with λ and \mathbf{f} standing for the eigenvalues and eigenvectors, respectively. In practice, we solve $\mathbf{P} \mathbf{f} = (1 - \lambda) \mathbf{f}$, where $\mathbf{P} = \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$ is symmetric, which corresponds to working with the normalized graph Laplacian. The coordinates \mathbf{Y} correspond to the $N \leq K - 1$ first eigenvectors associated to the N first eigenvalues sorted by ascending order after removing the trivial case associated to the eigenvalue zero.

In practice, we used the generalization of the manifold learning to DM (Coifman and Lafon, 2006), which provides invariance to the non-uniform density of the samples. We do so by first normalizing the affinity matrix \mathbf{W} by the total affinity of each sample to the other samples, namely using $\frac{W_{ij}}{D_i \cdot D_j}$ instead of W_{ij} . We used the simplest version of the diffusion time parameter, $t = 1$, which amounts at considering the diffusion between one sample and its neighbors in one step. Besides, the parameter that controls the influence of the density of samples was set to $\alpha = 1$, which corresponds to the Laplace-Beltrami approximation in the formalism of DM.

2.3.2. Two descriptors (fusion) - Multiple Kernel Learning

MKL is a manifold learning method that combines multiple descriptors at once and leads to a single low-dimensional space, contrary to MML. To do so, an affinity matrix is computed for each descriptor using an associated Gaussian kernel. MKL then combines the computed

affinity matrix and optimizes the following equation:

$$\Phi(\mathbf{Y}) = \sum_{i,j=1}^K \|\mathbf{A}^T \mathbb{K}^{(i)} \beta - \mathbf{A}^T \mathbb{K}^{(j)} \beta\|^2 \hat{W}_{ij} \quad (3)$$

where $\beta = [\beta_1, \dots, \beta_m]$ are the kernel weights with m the number of descriptors. \mathbf{A} is the matrix that maps the input into the low-dimensional space, \hat{W}_{ij} are the coefficients of a global affinity matrix and $\mathbb{K}^{(i)}$ is a matrix defined for the i -th sample by a combination of all the kernels from each descriptor (Lin et al., 2011). The MKL algorithm alternates the minimization of Eq. 3 with respect to either \mathbf{A} or β . The samples are then mapped to the latent space by: $\mathbf{Y} = \mathbf{A}^T \sum_{k=1}^m \mathbf{K}_k \beta_k$. This method will serve to benchmark our approach based on a alignment strategy (described in Section 2.3.4) against a fusion strategy.

2.3.3. Two descriptors (alignment) - Pairwise correspondences

(Ham et al., 2005) described a manifold learning method with a strict alignment strategy called pairwise correspondences. The method aims to link each sample of a latent space to the corresponding sample in the other latent space as defined by the following energy:

$$\Phi(\mathbf{Y}) = \sum_{i,j=1}^K \|\mathbf{y}_i^1 - \mathbf{y}_j^1\|^2 W_{ij}^1 + \sum_{i,j=1}^K \|\mathbf{y}_i^2 - \mathbf{y}_j^2\|^2 W_{ij}^2 + \mu \sum_{i=1}^K \|\mathbf{y}_i^1 - \mathbf{y}_i^2\|^2 \quad (4)$$

where $\mathbf{W}^1 = [W_{ij}^1]$ and $\mathbf{W}^2 = [W_{ij}^2]$ stand for the affinity matrices of the two descriptors and $\mu > 0$ is the weight of the interactions.

As for the diffusion maps approach, this problem can be reformulated using a matrix form, with the following affinity matrix:

$$\mathbb{W} = \begin{bmatrix} \mathbf{W}^1 & \mathbf{U}^{12} \\ \mathbf{U}^{21} & \mathbf{W}^2 \end{bmatrix} \quad (5)$$

where $\mathbf{U} = [U_{ij}]$ with $U_{ij} = \mu \mathbf{I}$.

With the matrix \mathbb{W} defined as above, Eq. 4 amounts at solving:

$$\Phi(\mathbf{Y}) = \text{tr}(\mathbf{Y}^T \mathbb{L} \mathbf{Y}) \quad (6)$$

$$\text{s.t. } \mathbf{Y}^T \mathbf{D}_{\mathbb{W}} \mathbf{Y} = \mathbf{I}$$

where $\mathbb{L} = \mathbf{D}_{\mathbb{W}} - \mathbb{W}$ and $\mathbf{D}_{\mathbb{W}} = \sum_j \mathbb{W}_{ij}$.

The minimization of Eq. 4 is achieved as for the single descriptor setting. As $\mathbb{W} \in \mathbb{R}^{2K \times 2K}$, we obtain eigenvectors whose first/last K rows correspond to the low-dimensional coordinates for the first/second descriptor respectively. The method therefore provides one latent space for each descriptor. Note that the formulation of Eq. 4 requires \mathbf{y}_i^1 and \mathbf{y}_i^2 to have the same dimensionality N .

2.3.4. Proposed approach: Two descriptors (alignment) - Multiple Manifold Learning

The pairwise correspondences method may be too strict to characterize cardiac shape and deformation, whose link may differ depending on the sample. MML is an alignment strategy that generalizes the pairwise correspondences method. MML weights the link between the descriptors through an inter-descriptors affinity matrix $\mathbf{M} = [M_{ij}] \in [0, 1]$, leading to the following energy to minimize:

$$\Phi(\mathbf{Y}) = \sum_{i,j=1}^K \|\mathbf{y}_i^1 - \mathbf{y}_j^1\|^2 W_{ij}^1 + \sum_{i,j=1}^K \|\mathbf{y}_i^2 - \mathbf{y}_j^2\|^2 W_{ij}^2 + \mu \sum_{i,j=1}^K \|\mathbf{y}_i^1 - \mathbf{y}_j^2\|^2 M_{ij} \quad (7)$$

where \mathbf{W}^1 and \mathbf{W}^2 stand for the affinity matrices of the two descriptors, and $\mathbf{M} = [M_{ij}] \in \mathbb{R}^{K \times K}$ encodes inter-descriptor correspondences, weighted by a factor $\mu > 0$.

In the way we defined \mathbb{W} in the pairwise correspondences, we can define the problem with the following block matrix:

$$\mathbb{W} = \begin{bmatrix} \mathbf{W}^1 & \mu \mathbf{M} \\ \mu \mathbf{M}^T & \mathbf{W}^2 \end{bmatrix} \quad (8)$$

We considered the formulation of the matrix \mathbf{M} (Valencia-Aguirre et al., 2011; Benkarim et al., 2020) where:

$$M_{ij} = \frac{\langle \mathbf{w}_i^1, \mathbf{w}_j^2 \rangle}{\|\mathbf{w}_i^1\| \|\mathbf{w}_j^2\|} \in [0, 1] \quad (9)$$

where \mathbf{w}_i^m is a row of the affinity matrix \mathbf{W}^m . Note that the values of \mathbf{w}_i^m lie in the interval $[0,1]$, and therefore $M_{ij} \in [0, 1]$.

This measure of similarity between descriptors is relatively simple, but interesting for our problem as it considers the relations between a sample and its nearest neighbors for each descriptor. The overall formulation of the problem would not differ with other similarity measures used in the literature, such as a generalization of the diffusion distance (Coifman and Hirn, 2014) or the wave kernel signature (Aubry et al., 2013; Clough et al., 2019).

In our implementation, the matrix \mathbf{M} by setting to 0 the elements of \mathbf{W}^1 and \mathbf{W}^2 that do not belong to the k_M nearest neighbors of each sample (i.e. the k_M highest value of each line of the matrix \mathbf{W}^1 and \mathbf{W}^2). \mathbf{M} is then computed by the Equation 9 with these two sparsified matrices. The solution of MML with this sparsification remains similar.

The minimization of Eq. 7 is achieved as for the pairwise correspondences setting. We describe in Supplementary Material how the formulation of the graph Laplacian associated to the matrix \mathbb{W} is obtained.

Of note, generalized versions of MML to more than two descriptors have been proposed in (Clough et al., 2019; Lindenbaum et al., 2020) which are addressed in a similar manner as above.

2.4. Exploiting the latent space

The following sections introduce one reconstruction method and two scalar measures that will be used in Section 3 to compare the latent spaces obtained by the different manifold learning methods.

2.4.1. Reconstruction from the low-dimensional spaces

Neither MML nor MKL have explicit operators to reconstruct high-dimensional samples from the low-dimensional space, which corresponds to the out-of-sample extension problem identified in (Bengio et al., 2004). Nonetheless, we can approximate this mapping by kernel regression. This regression looks for a mapping function f_m defined between the low- and high-dimensional spaces associated to the m -th descriptor, which minimizes:

$$\frac{1}{2} \|f_m\|^2 + \frac{\gamma_m}{2} \sum_{i=1}^K \|f_m(\mathbf{y}_i^m) - \mathbf{x}_i^m\|^2 \quad (10)$$

where γ_m is a trade-off between the regularization term (first term) and the adherence to the data (second term). The analytical solution of this problem is

$$f_m(\mathbf{y}^m) = \sum_{i=1}^K k(\mathbf{y}^m, \mathbf{y}_i^m) \mathbf{b}_i^m \quad (11)$$

where \mathbf{b}_i^m corresponds to the i -th column of the matrix:

$$\mathbf{B}^m = \left(\mathbf{K} + \frac{1}{\gamma_m} \mathbf{I} \right)^{-1} \mathbf{X}^m \quad (12)$$

with $\mathbf{K} = k(\mathbf{y}_i^m, \mathbf{y}_j^m) \in \mathbb{R}^{K \times K}$ and \mathbf{I} the identity matrix. The kernel function k is defined as a Gaussian with imposed bandwidth σ . In practice, we used a multiscale implementation (Duchateau et al., 2013), which is more robust to the non-uniform density of the samples and only requires to set the hyperparameter γ_m .

2.4.2. Neighborhood consistency

The closest neighbors to each sample are easily identifiable from the affinity matrix associated to each descriptor (for MML and DM). We exploited this information to compute the Euclidean distance in each latent space between each sample and its n nearest neighbors as:

$$d_i = \frac{1}{n} \sum_{j \in V_i} \|\mathbf{y}_i^m - \mathbf{y}_j^m\|^2 \quad (13)$$

with V_i representing the n nearest neighbors to the sample i in the input space. Before computing the distance, each latent space was normalized according to the standard deviation along the first dimension.

2.4.3. Anomaly quantification

Shape and deformation anomalies were quantified from the latent spaces by comparing individual samples to the controls, used as reference for normality. This can be achieved through a leave-one-out strategy: for each sample, the MML and MKL latent spaces are computed from the remaining $K-1$ samples, and the excluded sample is projected onto these latent spaces. The projection was performed by kernel regression for MML (Duchateau et al., 2013), and the intrinsic projection provided by MKL as in (Sanchez-Martinez et al., 2017). Finally, anomalies were quantified in the latent space between the

tested sample and the distribution of controls (reference for normality) by the Mahalanobis distance, which is relevant to quantify the distance between one sample and a distribution. The process was repeated for all of the samples in the population.

3. Experiments

Our Matlab code for DM and MML is publicly available¹, accompanied with demo data corresponding to the synthetic experiment discussed in this section. The Matlab code for MKL corresponds to the one from Sanchez-Martinez et al. (2017).

3.1. Experiments on synthetic data

The four approaches (DM, MML, MKL and *pairwise correspondences*) require defining the width σ of the Gaussian kernel, which we set as the average distance of each point to its k_σ -th neighbor. MML is also controlled by μ (the weight of the interactions) and k_M (the amount of nearest neighbors to sparsify the extra-diagonal matrix \mathbf{M}).

3.1.1. Choice of hyperparameters

To illustrate the choice of the hyperparameters for MML, we specifically designed a synthetic database where the link between the descriptors is fully controlled. A OD model was used to simulate pressure and volume curves across one cycle of the left ventricle. The model is written in CellML² and can be launched in the OpenCor modelling platform (Molléro et al., 2018). We used a script version provided by the authors that facilitates repeated runnings of the model with different sets of parameters. A first set of 130 cases was obtained by only modifying two parameters among 14 active parameters: the relaxation rate k_{rs} (in the range 5 s^{-1} to 60 s^{-1}) and the maximal contraction σ_0 (in the range 4 MPa to 10 MPa). We selected these two parameters because of their influence on the pressure and volume (see Fig. 2 with varying σ_0 and fixed $k_{rs} = 10 \text{ s}^{-1}$). In this configuration, pressure and volume are outputs of the same model and are therefore fully linked.




Figure 2: Simulation of the left ventricular pressure (left) and volume (right) across a full cardiac cycle, for a fixed $k_{rs} = 10 \text{ s}^{-1}$ and a varying σ_0 (MPa).

This link was deteriorated by changing a third parameter (the peripheral resistance R_p) for one sample and its 8 neighbors regarding σ_0 and k_{rs} (σ_0 equal to 6.5, 7.0 or 7.5 MPa and k_{rs} equal to 16, 20 and 25 s^{-1}). For those samples, R_p (MPa.m⁻³.s) was selected randomly in a bandwidth centered around the initial value $R_p = 2.5 \times 10^{-7} \text{ MPa.m}^{-3}.\text{s}$.

The influence of the hyperparameters μ and k_M on the latent space is illustrated in Fig. 3. We also tested several values of k_σ around 10% of the number of samples, but these had little influence on the latent space and are not illustrated here. We set $k_\sigma = 10$ for the experiments on this dataset. Rows respectively correspond to the volume and pressure low-dimensional spaces and the point-to-point differences between them. The nine points for which the pressure-volume link was artificially deteriorated are highlighted in red in the last row. For the volume and pressure latent spaces, the samples are colored by their value of k_{rs} and σ_0 .

Due to the non-linear changes observed in the curves depicted in Fig. 2, we only compared non-linear methods in this synthetic experiment. DM consider the descriptors independently, and substantial differences between the two latent spaces are observed (last row of Fig. 3a), even if we applied an a-posteriori Procrustes alignment. For MML, as the descriptors are fully correlated except for nine noisy points, we expect the two latent spaces to be close on the last row. Despite some noisy points, the latent spaces are consistent as assessed by the arrangement of the colored samples.

At fixed k_M (Fig. 3b), the optimal μ was obtained by computing the energy in Eq. 7 for several values of μ at

¹https://gitlab.com/maxDif/mml_pressure-volume-demo

²<https://models.cellml.org/workspace/44c>

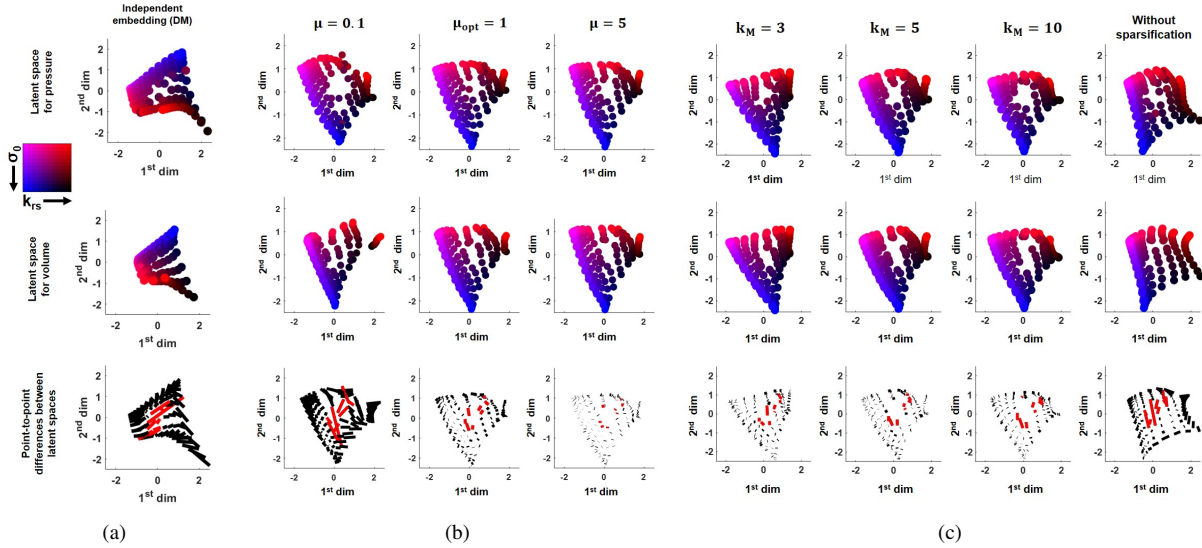


Figure 3: Influence of the hyperparameters μ and k_M , tested on the synthetic population (a) Independent embedding using DM after Procrustes alignment. (b) Latent spaces using MML for a fixed $k_M = 10$ and a varying μ . (c) Latent spaces using MML without sparsification and with different k_M at optimal μ ($\mu = 0.9, 1, 1$ and 0.1 respectively for $k_M = 3, k_M = 5, k_M = 10$ and without sparsification). The rows from top to bottom respectively represent the latent spaces for pressure, volume and the difference between them. The samples are colored by their k_{rs} and σ_0 values, as summarized in the small image on the left. The red lines in the bottom row plots correspond to samples whose link between the two descriptors was artificially deteriorated.

fixed $k_{\sigma} = 10$ and $k_M = 10$, and picking the value at which it is minimized. We observe that the higher μ , the lower the differences between the latent spaces. This confirms that the hyperparameter μ gives more weight to the link between the descriptors and forces the latent spaces to be closer. Nevertheless, a high μ enforces the link and is not optimal regarding the energy defined in Eq. 7.

We also analyzed the influence of k_M when μ is optimal (Fig. 3c). We applied MML without and with sparsification of \mathbf{M} and for different values of k_M . As in the previous experiment, the latent spaces are coherent because they are arranged according to their values of k_{rs} and σ_0 . As visible in the first three columns, differences between the latent spaces are low and mainly affect the noisy points. Without sparsification of k_M (the last column), the difference increases and all points are affected.

Based on these experiments, we would opt for this dataset for $k_M = 10$, to avoid substantial influence of the noisy points, before selecting the other hyperparameters (for this dataset $k_{\sigma} = 10$ and $\mu = 1$).

3.1.2. Comparison between MML and pairwise correspondences

We also compared MML and the *pairwise correspondence* with the stricter formulation from (Ham et al., 2005), which forces pairwise correspondences between samples. Figure 4 shows the results obtained with the optimal hyperparameters for MML and Ham’s method. We can observe that the *pairwise correspondence* formulation fails to apprehend the link between the descriptors, compared to MML which succeeds to keep close the points where the link is not deteriorated. This behaviour is due to the matrix \mathbf{M} in MML. Indeed, if samples i (descriptor 1) and j (descriptor 2) are strongly linked, meaning according to Eq. 9 that their neighborhood is comparable, then M_{ij} is close to 1, and the MML cost function attempts to bring closer the coordinates \mathbf{y}_i^1 and \mathbf{y}_j^2 . Conversely, if M_{ij} is close to 0, the coordinates \mathbf{y}_i^1 and \mathbf{y}_j^2 will not necessarily be brought close to each other. In contrast, the *pairwise correspondence* scheme forces this matching, but only for the pairs \mathbf{y}_i^1 and \mathbf{y}_i^2 (and not \mathbf{y}_j^2).

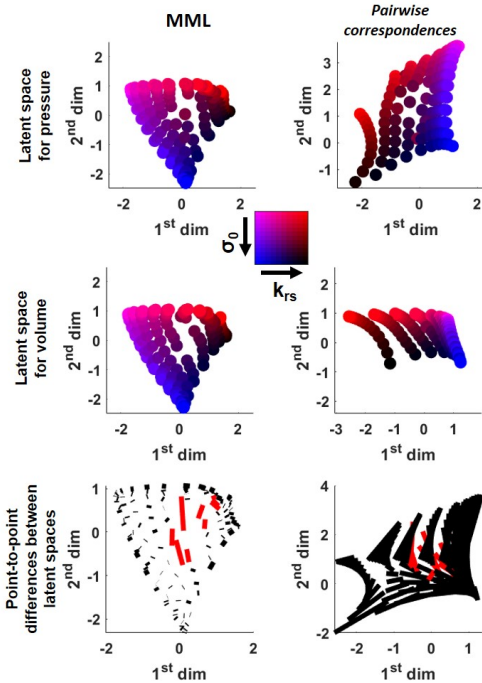


Figure 4: Comparison of the latent spaces obtained with MML and the pairwise correspondences scheme, each with optimal hyperparameters. The rows from top to bottom respectively represent the latent spaces for pressure, volume and the difference between them. The samples are colored by their k_{rs} and σ_0 values, as in Fig. 3

3.2. Experiments on real cases with RV pressure overload

In this section, we applied our methodology to a subset of the RV database consisting of 150 PH patients and 100 controls. Five patients were excluded due to obvious mesh defects near the valves that produced very localized strain artifacts, and resulted in strong outliers that incorrectly constrained the first embedding dimensions. Lower deformation and more dilated shapes have been reported for worsening grades of PH (Moceri et al., 2018) using a pixel-wise analysis of deformation patterns displayed on subgroup shapes. We propose to confront these findings with our methodology, which examines more in detail the population distribution at the pattern level. The hyperparameters were determined using an approach similar to the one described in the previous section. On this real dataset, the retained values were: $k_\sigma = 10$ for all the meth-

ods and $k_M = 10$ and $\mu = 1$ for MML.

3.2.1. Latent spaces

Figure 5 represents the latent spaces obtained with DM, MML and MKL applied to both descriptors (DM was applied to each descriptor independently).

The dimensions from the shape and deformation latent spaces are weakly correlated for DM ($r^2 = 0.07, 0.08$ and 0.06 for the 1st, 2nd and 3rd dimensions respectively), unlike for MML ($r^2 = 0.84, 0.84$ and 0.85 , respectively). Besides, the DM latent space for deformation exhibits strong outliers, while MML and MKL appear more robust to this effect.

The clinical objective is not to classify the subjects by their subgroup, because the diagnostic is already known, but to apprehend the variability inside the population. Nevertheless, the subgroup arrangement for healthy and PH is a good indicator for the coherence of our results.

In particular, we observe that the first dimension of the deformation latent spaces roughly separates the two subgroups. This dimension is associated to the strain amplitude as illustrated in Section 3.2.2.

3.2.2. Latent dimensions

We also explored the main variations encoded in the latent spaces. For each strategy, we reconstructed the high-dimensional data (shape and deformation) for samples at $-2\sigma, 0$, and $+2\sigma$ as displayed in Fig. 6. Animations of the latent dimensions are available as supplementary movies. Contrary to MKL which provides a single space, this process was performed on each latent space independently for DM and MML, and combined into a single subplot for the sake of simplicity.

The average shape and deformation patterns are very similar across methods, as confirmed in the difference plotted in Fig. 7. Strain differences are lower than 10% between MML and DM, MML and MKL, and are mainly observed around the valve which is a region where the tracking and remeshing from the commercial software is less reproducible. Shape differences are very low, the maximal distance being 1 mm (of note, the shapes height is around 80-90mm from base to apex).

The methods differ more regarding the principal latent dimensions, as shown in Fig. 6. With MML, the first

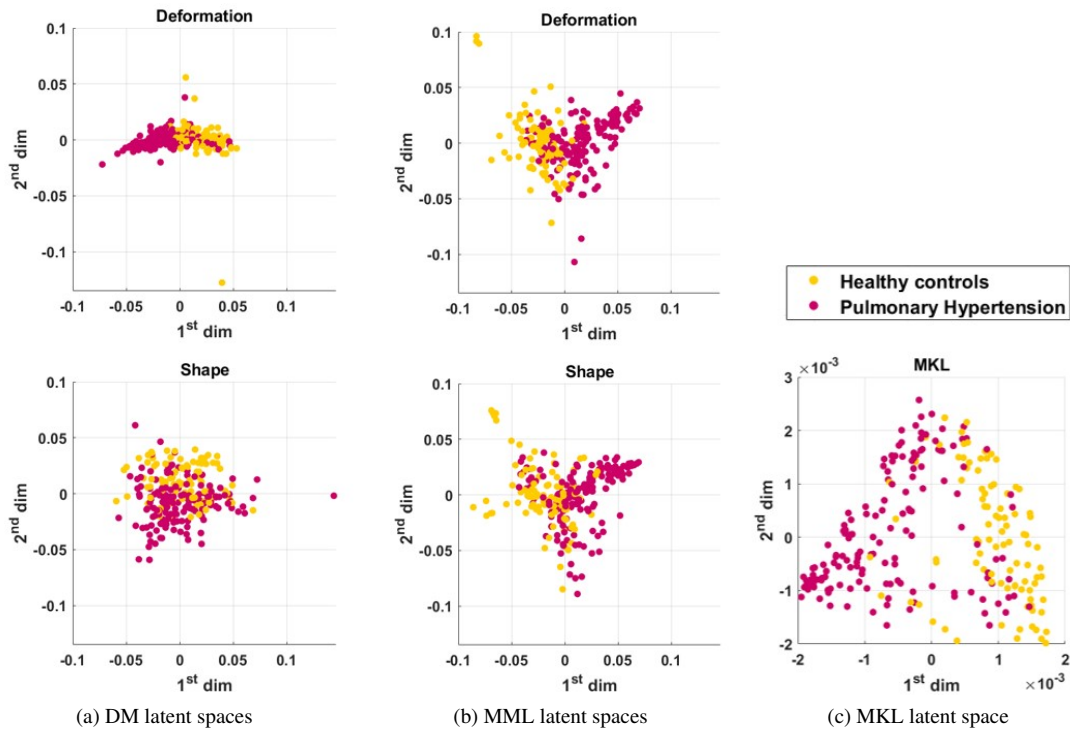


Figure 5: Latent spaces provided by DM, MML and MKL for the controls and pulmonary hypertension subsets of the RV population colored by their labels.

dimension encodes both shape and deformation variations that are characteristic of PH and control differences: more spherical shapes and in particular higher curvature of the septum (black arrow), and lower deformation are observed for more severe PH patients. The second dimension encodes shape differences that can be found on both subgroups, in particular the distance between the valves. Similar observations can be made for MKL. In contrast, DM does not capture the link between shape and deformation. The strain patterns are comparable to those found for MML and MKL along the first dimension, but very different patterns are observed along the second dimension. Similarly, the bowing of the septum, which is characteristic of the studied disease, is more visible on the second dimension, while the intra-valve distance dominates the first dimension.

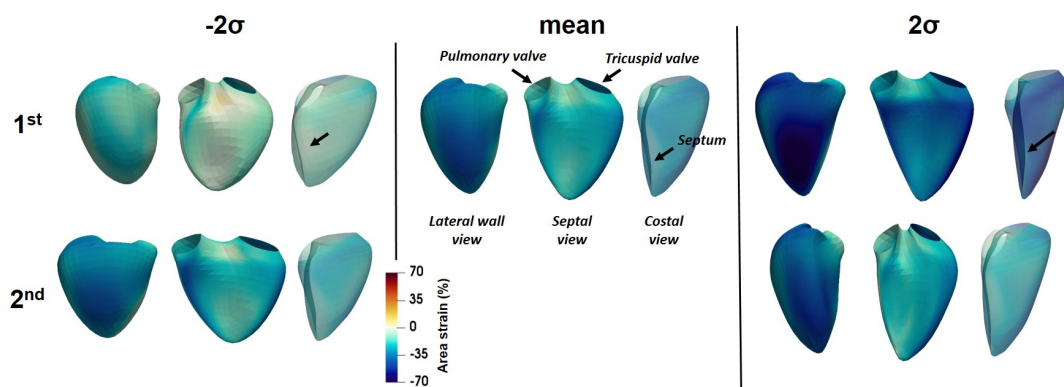
In order to confirm these observations quantitatively, we computed global indicators (global area strain, and

volume) along the latent dimensions in Fig. 8. The global area strain is similar between all the methods. Concerning volume, as observed in Fig. 6, the second mode of DM is similar to the first one for MML and MKL. We also notice that the second latent dimension of MML and MKL is almost invariant to the volume.

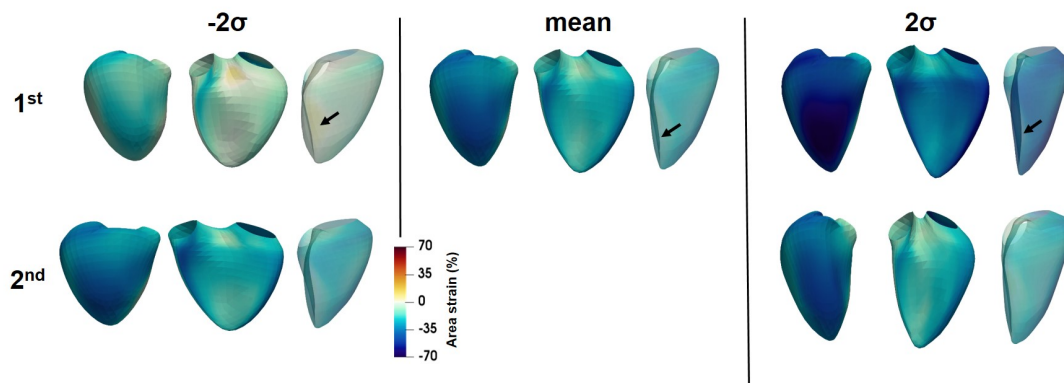
These results confirm the relevance of considering both descriptors at the same time, although MML and MKL seem to have similar behaviors. The following section will evaluate the relevance of two vs. one latent space respectively provided by MML and MKL.

3.2.3. Neighborhood consistency

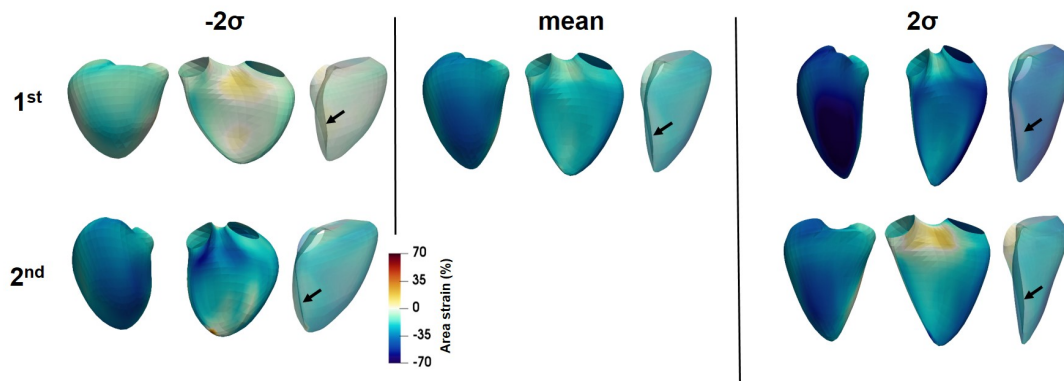
In this section, we analyzed further the results of Section 3.2.1 and shown in Fig. 5. We used the distance described in Section 2.4.2 and generalized this experiment to all samples. The distribution of the neighbors distance for MML, DM and MKL and each possible affinity ma-



(a) MML latent dimensions



(b) MKL latent dimensions



(c) DM latent dimensions

Figure 6: The first two latent dimensions for PH and control subsets, obtained with MML, MKL and DM. For better visualization of shape differences, the lateral wall and septal views are complemented by a third costal view. Transparency was set to 0.5 on this view to better appreciate fine changes in the septal curvature, somehow masked by the crescent shape of the RV.

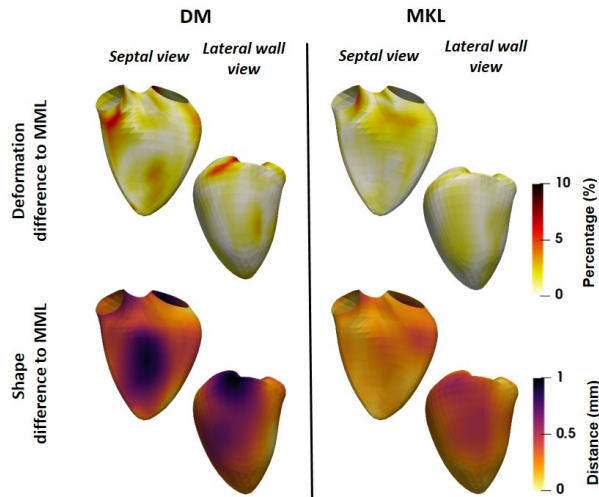


Figure 7: Average shape and deformation differences between DM and MKL, each compared to the MML average of the PH and controls subsets.

trix (deformation or shape) is shown in Fig. 9. For the deformation and shape neighbors, the neighborhoods are better preserved with DM and MML, compared to MKL. DM present outliers that correspond to the outlying samples visible in the latent space. MML succeeds to keep the neighbors consistency close to DM for each descriptor, contrary to MKL which fuses the descriptors into one single representation and deteriorates it (higher distribution of the mean neighbors distance).

3.2.4. Comparison to linear embedding by PLS

In this section, we compared the MML results to the ones obtained with PLS, a linear approach used in our previous work (Di Folco et al., 2020), which involves computing the cross-covariance between the two descriptors. A similarity transform was applied for shape normalization to remain consistent with what was used for MML (a Procrustes alignment was used in (Di Folco et al., 2020)). Figure 10a reveals that PLS provides slightly different average shape and deformation from those obtained with MML, with differences comparable to those noted for DM and MKL in Fig. 7. Despite similarities regarding the average, PLS may reconstruct some non-plausible shapes far from the average along the 2nd and 3rd latent dimen-

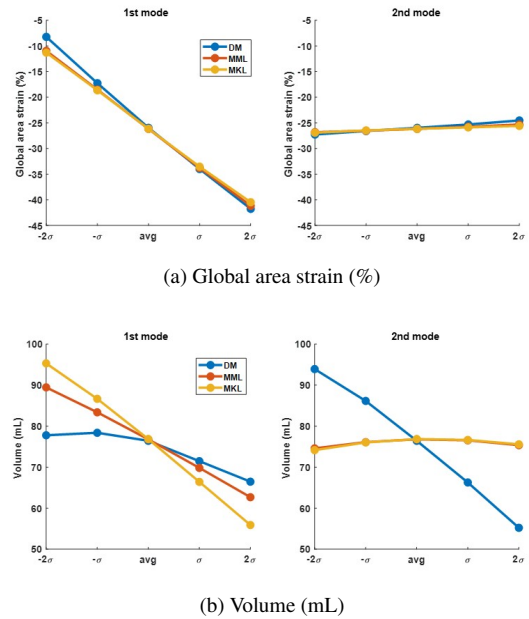


Figure 8: Variations of global indicators (global area strain and volume) along the first two latent dimensions for DM, MML and MKL applied to the control and PH population.

sions (valves too close to each other, and abnormal membranous septum between the two valves, respectively in Fig. 10b).

3.3. Experiments on real cases with RV volume overload

This section focuses on another subset of the RV database, which consists of the 27 ASD and the 28 ToF patients, and their 55 corresponding age- and sex-matched controls. Using global scalar values (volume and global strain (Moceri et al., 2018)) or average shapes and deformation patterns (Moceri et al., 2020), the literature reported that ASD and ToF have dilated shapes, but ASD has preserved deformation compared to ToF. We investigated whether these observations are confirmed by quantifying anomalies using the latent spaces provided by MML and MKL and a careful examination of the latent spaces as described in Section 2.4.3. The retained hyperparameters values for this dataset were: $k_{\sigma} = 10$ for all the methods, and $k_M = 10$ and $\mu = 1$ for MML.




Figure 9: Distribution of the mean neighbors distance for DM, MML and MKL applied to the control and PH population for the deformation and shape descriptors.


3.3.1. Latent spaces

Fig. 11 shows the latent spaces of the population for MML and MKL. The first two dimensions of the deformation and shape latent spaces for MML are highly correlated ($r^2 = 0.82$ and $r^2 = 0.80$ for the 1st and 2nd respectively) but are not correlated to the MKL latent space (Deformation/MKL: $r^2 = 0.22$ and Shape/MKL: $r^2 = 0.29$). We can observe that the deformation latent space separates the healthy controls from the ToF subjects. In contrast, ASD and ToF are closer in the MML shape latent space and in the MKL latent space, and rather separated from the healthy subjects.

3.3.2. Anomaly quantification

To analyze further the latent spaces, we quantified anomalies using the distance described in Section 2.4.3 for the different subgroups.

As shown in Fig. 12, the three subgroups exhibit similar anomalies using the single latent space provided by MKL. In contrast, the separated latent spaces provided by MML confirm the observations made in the literature. Regarding deformation, ASD patients have low anomalies comparable to the controls, unlike ToF patients. Regarding shape, ASD and ToF patients presented higher anomalies as compared to controls. Of note, controls are less consistent regarding shape compared to deformation. This experiment confirms the benefit of having separate but aligned latent spaces to better represent each descriptor.



(a) Difference to the MML average (b) Sample of latent dimensions

Figure 10: Comparison between MML and the PLS linear approach for the PH population. (a) Shape and deformation differences between PLS and MML averages, (b) Unplausible shapes observed along the 2nd and 3rd latent dimensions for PLS, compared to MML.

4. Discussion

In this paper, we proposed a strategy to characterize the interactions between two partially related physiological traits of the cardiac function: myocardial shape and deformation, assessed through high-dimensional descriptors defined at each point of the myocardium. We exploited a non-linear manifold learning method for manifold alignment (MML) to both control the representation of individual descriptors and the relationship between two different descriptors. We demonstrated its value for the characterization of several RV diseases where myocardial shape and deformation are partially related (depending on the disease type and grade).

We benchmarked the MML approach against the independent embedding of each descriptor (using DM), and a fusion strategy that provides a single latent space (MKL). We observed some similarities between the low-dimensional latent spaces. However, a finer analysis of the latent dimensions showed that a joint approach (fusion or alignment) better captured the link between shape and deformation variations compared to a strategy that considers them independently. Furthermore, several experiments on the neighborhood consistency and the anomaly




Figure 11: Latent spaces for the controls, ASD and ToF subsets of the RV population provided by MML and MKL colored by their labels.




Figure 12: Distribution of anomalies across the ASD, ToF and control subjects, quantified by the Mahalanobis distance to the controls, from the two latent spaces provided by MML and the single latent space provided by MKL.

quantification demonstrated the assets of manifold alignment, which provides related spaces without potential loss of information, compared to a fusion strategy. This is crucial for our application, given that the link between the two descriptors is unknown. Moreover, we related the main variations and the anomaly quantification experiment to clinical observations that confirmed the pertinence of the approach.

We tested beforehand our method on a synthetic dataset where the link is fully controlled, to apprehend the differ-

ent hyperparameters and notably the amount of sparsification of the affinity matrix controlling the interactions between descriptors. We demonstrated the robustness of the sparsification strategy to a deteriorated link between descriptors. The notable influence of the hyperparameter μ was also investigated. A reasoning to select the hyperparameters based on our observations was presented: we first selected k_M (to limit the influence of low correlated samples), then k_σ and finally the optimal μ .

MML was applied in literature work to align datasets belonging to the same manifold with either different images of objects undergoing the same orientation changes (Valencia-Aguirre et al., 2011), or 2D MRI slices at different positions (Baumgartner et al., 2017; Clough et al., 2019). In our work, we aimed at explicitly characterizing the link between two physiologically related descriptors, namely cardiac shape and deformation, using the MML framework. However, these two descriptors do not necessarily correspond to the same manifold, as their link may vary across individuals and disease.

Besides, we have shown that the MML strategy is interesting for this purpose compared to the stricter alignment from (Ham et al., 2005), which enforces pairwise correspondences for all samples. Our experiments on synthetic data confirmed that MML is able to leverage the strength of the alignment depending on the actual link between the two descriptors, which is key for our application. This setting has not been explored before in the literature, which considered multiple datasets originating from a single manifold.

The MML algorithm involves an affinity matrix expressing the relations between samples and descriptors, whose definition naturally conditions the embedding. We opted for a simple formulation of the inter-descriptor distances to better focus on the control of the descriptors interactions. We sparsified the matrix encoding such interactions to reinforce the link between one sample and its closest neighbors. Stronger sparsification schemes up to one-to-one correspondences have been used (Clough et al., 2019), but these do not suit our problem where considering several neighbors should be preferred for each sample, as illustrated on our synthetic experiment.

In our preliminary work (Di Folco et al., 2020), we related shape and deformation using PLS, a linear dimensionality reduction method that considers the cross-covariance between the two descriptors. This work re-

vealed that considering the two descriptors jointly or independently leads to substantial differences in areas of critical importance for the studied diseases. However, linear methods may be limited for shape and deformation, which most-probably originate from a non-linear manifold. Although this limitation may be subtle to assess, we pointed out that PLS can produce some unplausible shapes when moving far from the average along some latent dimensions. Note that PLS may be harder to extend to multiple descriptors and does not explicitly control the weight given to the interactions, contrary to MML (Clough et al., 2019).

Linear methods such as PLS come with an explicit reconstruction method. For non-linear manifold learning, reconstruction from the latent spaces is done a-posteriori and requires approximation operators. For MML, we used kernel regression, which might have softened some potential inconsistencies for extreme cases. This aspect may be improved by the explicit encoding/decoding of variational autoencoders, but they currently provide less flexibility on the statistical representation of populations.

The Procrustes alignment used a similarity transform and therefore removed scaling differences between subjects, which differs from our previous work (Di Folco et al., 2020) where we used a rigid transform. This allows observing finer shape differences compared to global differences previously reported in the literature (Dragulescu et al., 2013). Besides, with rigid transformations, very small or large shapes in our database are considered outliers due to insufficiently close neighbors, and these characteristics dominate the first dimensions of the latent space. This might be a limitation of manifold learning methods that may require lots of data in specific regions of the manifold. This effect can be attenuated by the use of Diffusion Maps, which can provide invariance to the non-uniform density of the samples, and by using a generalized Procrustes with a similarity transformation as we did. The clinical literature often reports normalizations by the body mass index, or more specific ones such as the body size (Medrano-Gracia et al., 2014). The proper normalization of shape differences and their effect on functional descriptors such as strain is a topic under investigation (Guigui et al., 2019, 2021) and the object of our future development of the present work.

Finally, our work does not consider confounding variables such as QRS duration, RV wall thickness or the

strong interactions between the Left Ventricle (LV) and the RV, which can affect the relationship between shape, deformation, and disease. Nonetheless, we could demonstrate the value of jointly considering shape and deformation on different RV diseases where such descriptors are partially related. Our observations on PH patients highlight key characteristics of the disease that are missed by an independent analysis (septal bowing, accompanied by more global shape dilatation and reduced deformation (Moceri et al., 2018)). Our work confirms the differences between ASD and ToF patients previously reported with single scalar values (Moceri et al., 2020). Moreover, our analysis goes further by examining high-dimensional descriptors and the distribution of samples, not limited to average trends (Moceri et al., 2020), hence paving the ground for a finer analysis of disease occurrence and evolution in such populations.

5. Conclusion

We presented a strategy to characterize the interactions between cardiac shape and deformation using non-linear manifold learning. We built upon an alignment approach with multiple manifold learning and evaluated the relevance of joint embedding on a synthetic example and real data from a RV imaging cohort. Our method was benchmarked against a fusion approach (MKL) and an approach that considers descriptors independently (DM). Results obtained with MML showed better coherence between the latent spaces associated to each descriptor, and an increased robustness to outliers. For the application we target, we also demonstrated the benefit of having several linked latent spaces as provided by alignment methods, compared to a single one provided by fusion approaches. Such latent spaces confirmed in a more detailed manner the clinical observations reported for the RV diseases studied here.

Acknowledgments

The authors acknowledge the partial support from the French ANR (MIC-MAC JCJC project [ANR-19-CE45-0005], and LABEX PRIMES of Université de Lyon [ANR-11-LABX-0063] within the program Investissements d’Avenir [ANR-11-IDEX-0007]), and the EEA

doctoral school. They also thank Gabriel Bernardino (CREATIS) for insightful discussions.

References

- Aubry, M., Schlickewei, U., Cremers, D., 2013. The wave kernel signature: A quantum mechanical approach to shape analysis. 2011 IEEE ICCV Workshops , 1626–33doi:10.1109/ICCVW.2011.6130444.
- Baumgartner, C.F., Gomez, A., Koch, L.M., Housden, J.R., Kolbitsch, C., McClelland, J.R.o., 2015. Self-Aligning Manifolds for Matching Disparate Medical Image Datasets. Proc. IPMI, LNCS 24, 363–374. doi:10.1007/978-3-319-19992-4_28.
- Baumgartner, C.F., Kolbitsch, C., McClelland, J.R., Rueckert, D., King, A.P., 2017. Autoadaptive motion modelling for MR-based respiratory motion estimation. Med Image Anal 35, 83–100. doi:10.1016/j.media.2016.06.005.
- Bengio, Y., Paiement, J.F., Vincent, P., Delalleau, N., Le Roux, N., Ouimet, M., 2004. Out-of-sample extensions for LLE, Isomap, MDS, Eigenmaps, and Spectral Clustering. Proc. NIPS , 177–84.
- Benkarim, O., Piella, G., Rekik, I., Hahner, N., Eixarch, E., Shen, D., et al., 2020. A novel approach to multiple anatomical shape analysis: Application to fetal ventriculomegaly. Med Image Anal 64, 101750. doi:10.1016/j.media.2020.101750.
- Bijnens, B., Cikes, M., Butakoff, C., Sitges, M., Crispi, F., 2012. Myocardial motion and deformation: What does it tell us and how does it relate to function? Fetal Diagn Ther 32, 5–16. doi:10.1159/000335649.
- Cikes, M., Solomon, S.D., 2016. Beyond ejection fraction: an integrative approach for assessment of cardiac structure and function in heart failure. Eur Heart J , 1642–50doi:10.1093/eurheartj/ehv510.
- Clough, J.R., Balfour, D., Cruz, G., Marsden, P.K., Prieto, C., Reader, A.J., et al., 2019. Weighted manifold alignment using wave kernel signatures for aligning medical image datasets. IEEE Trans Pattern Anal Mach Intell 42, 988–97. doi:10.1109/TPAMI.2019.2891600.
- Coifman, R.R., Hirn, M.J., 2014. Diffusion maps for changing data. Applied Comput Harm Anal 36, 79–107. doi:10.1016/j.acha.2013.03.001.
- Coifman, R.R., Lafon, S., 2006. Diffusion maps. Applied Comput Harm Anal 21, 5–30. doi:10.1016/j.acha.2006.04.006.
- Coifman, R.R., Lafon, S., Lee, A.B., Maggioni, M., Nadler, B., Warner, F., et al., 2005. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. Proc Natl Acad Sci 102, 7426–31. doi:10.1073/pnas.0500334102.
- Di Folco, M., Clarysse, P., Mocerri, P., Duchateau, N., 2020. Learning interactions between cardiac shape and deformation: Application to pulmonary hypertension Proc. STACOM-MICCAI’19, LNCS, 119–27. doi:10.1007/978-3-030-39074-7_13.
- Dragulescu, A., Grosse-Wortmann, L., Redington, A., Friedberg, M.K., Mertens, L., 2013. Differential effect of right ventricular dilatation on myocardial deformation in patients with atrial septal defects and patients after tetralogy of Fallot repair. Int J Cardiol 168, 803–810. doi:10.1016/j.ijcard.2012.10.009.
- Duchateau, N., De Craene, M., Piella, G., et al., 2012. Constrained manifold learning for the characterization of pathological deviations from normality. Med Image Anal 16, 1532–49. doi:10.1016/j.media.2012.07.003.
- Duchateau, N., De Craene, M., Sitges, M., Caselles, V., 2013. Adaptation of multiscale function extension to inexact matching: Application to the mapping of individuals to a learnt manifold. Proc. SEE-GSI, LNCS 8085, 578–86. doi:10.1007/978-3-642-40020-9_64.
- Duchateau, N., King, A.P., De Craene, M., 2020. Machine learning approaches for myocardial motion and deformation analysis. Front Cardiovasc Med 6, 190. doi:10.3389/fcvm.2019.00190.
- Gilbert, K., Mauger, C., Young, A.A., Suinesiaputra, A., 2020. Artificial intelligence in cardiac imaging with statistical atlases of cardiac anatomy. Front Cardiovasc Med 7, 102. doi:10.3389/fcvm.2020.00102.

- Gower, J.C., 1975. Generalized Procrustes analysis. *Psychometrika* 40, 33–51. doi:10.1007/BF02291478.
- Guigui, N., Jia, S., Sermesant, M., Pennec, X., 2019. Symmetric algorithmic components for shape analysis with diffeomorphisms. *Proc. GSI, LNCS 11712*, 759–68. doi:10.1007/978-3-030-26980-7_79.
- Guigui, N., Mocerri, P., Sermesant, M., et al., 2021. Cardiac motion modeling with parallel transport and shape splines. *Proc. ISBI*, 1394–1397.
- Ham, J., Lee, D.D., Saul, L.K., 2005. Semisupervised alignment of manifolds. *Proc. AISTATS 120*, 120–7.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A., 2016. beta-vaes: Learning basic visual concepts with a constrained variational framework. *OpenReview*.
- Kingma, D.P., Welling, M., 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Lawrence, N.D., Moore, A.J., 2007. Hierarchical gaussian process latent variable models. *Proc. ICML*, 481–8.
- Lee, C.S., Elgammal, A., Torki, M., 2016. Learning representations from multiple manifolds. *Pattern Recognit* 50, 74–87. doi:10.1016/j.patcog.2015.08.024.
- Li, Y., Yang, M., Zhang, Z., 2018. A survey of multi-view representation learning. *IEEE Trans Knowl Data Eng* 31, 1863–83. doi:10.1109/TKDE.2018.2872063.
- Lin, Y., Liu, T., Fuh, C.S., 2011. Multiple kernel learning for dimensionality reduction. *IEEE Trans Pattern Anal Mach Intell* 33, 1147–60. doi:10.1109/TPAMI.2010.183.
- Lindenbaum, O., Yeredor, A., Salhov, M., Averbuch, A., 2020. Multi-view diffusion maps. *Inform Fusion* 55, 127–49. doi:10.1016/j.inffus.2019.08.005.
- Medrano-Gracia, P., Cowan, B.R., Ambale-Venkatesh, B., Bluemke, D.A., Eng, J., et al., 2014. Left ventricular shape variation in asymptomatic populations: the Multi-Ethnic Study of Atherosclerosis. *Journal of Cardiovascular Magnetic Resonance* 16, 56.
- Mocerri, P., Duchateau, N., Baudouy, D., Schouver, E.D., Leroy, S., Squara, F., et al., 2018. Three-dimensional right-ventricular regional deformation and survival in pulmonary hypertension. *Eur Heart J Cardiovasc Imaging* 19, 450–8. doi:10.1093/ehjci/jex163.
- Mocerri, P., Duchateau, N., Gillon, S., Jaunay, L., Baudouy, D., Squara, F., et al., 2020. Three-dimensional right ventricular shape and strain in congenital heart disease patients with right ventricular chronic volume loading. *Eur Heart J Cardiovasc Imaging* doi:10.1093/ehjci/jeaa189. in press.
- Molléro, R., Pennec, X., Delingette, H., Garny, A., Ayache, N., Sermesant, M., 2018. Multifidelity-CMA: a multifidelity approach for efficient personalisation of 3D cardiac electromechanical models. *Biomech Model Mechanobiol* 17, 285–300. doi:10.1007/s10237-017-0960-0.
- Puyol-Antón, E., Sinclair, M., Gerber, B., Amzulescu, M.S., Langet, H., De Craene, M., et al., 2017. A multimodal spatiotemporal cardiac motion atlas from MR and ultrasound data. *Med Image Anal* 40, 96–110. doi:10.1016/j.media.2017.06.002.
- Sanchez-Martinez, S., Duchateau, N., Erdei, T., Fraser, A.G., Bijnens, B.H., Piella, G., 2017. Characterization of myocardial motion patterns by unsupervised multiple kernel learning. *Med Image Anal* 35, 70–82. doi:10.1016/j.media.2016.06.007.
- Sanz, J., Sánchez-Quintana, D., Bossone, E., Bogaard, H.J., Naeije, R., 2019. Anatomy, function, and dysfunction of the right ventricle. *J Am Coll Cardiol* 73, 1463–82. doi:10.1016/j.jacc.2018.12.076.
- Tenenbaum, J.B., Silva, V., Langford, J.C., 2000. A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 2319–23. doi:10.1126/science.290.5500.2319.
- Valencia-Aguirre, J., Álvarez Meza, A., Daza-Santacoloma, G., Acosta-Medina, C., Castellanos-Domínguez, C.G., 2011. Multiple manifold learning by nonlinear dimensionality reduction. *Proc. CIARP, LNCS 7042*, 206–13. doi:10.1007/978-3-642-25085-9_24.

Wang, B., Mezlini, A.M., Demir, F., Fiume, M., Tu, Z., et al., 2014. Similarity network fusion for aggregating data types on a genomic scale. *Nat Methods* 11, 333–337. doi:10.1038/nmeth.2810.

Yan, S., Xu, D., Zhang, B., Zhang, H., Yang, Q., Lin, S., 2007. Graph embedding and extensions: A general framework for dimensionality reduction. *IEEE Trans Pattern Anal Mach Intell* 29, 40–51. doi:10.1109/TPAMI.2007.250598.