

Visualizing definitional divergence in high-dimensional data by manifold alignment: Application to 3D right ventricular strain computations

Maxime Di Folco, Gabriel Bernardino, Patrick Clarysse, and Nicolas Duchateau

Abstract—Medical imaging studies often rely on a single sample per subject, assuming it is representative of their physiological traits. However, variations in how input descriptors are defined or computed (e.g. due to a lack of consensus in the scientific field) may have a crucial impact on the analysis, and are hardly considered in practice. In this paper, we propose an original strategy based on representation learning to estimate a parametric map reflecting the impact of such definitional differences on a given physiological descriptor, previously extracted from medical images. We consider the different definitions or computations of such physiological descriptors as different high-dimensional data, potentially of heterogeneous types. We specifically focus on myocardial deformation (strain), for which there is limited agreement on its definition. We first use manifold alignment to match the latent representations associated with the different definitions of this descriptor. Then, we formulate plausible distributions in the latent space to represent definitional divergence across descriptors, from which we reconstruct a high-dimensional parametric map to visualize such definitional divergence.

Due to the lack of proper ground truth for this specific clinical application, we first demonstrate this methodology on toy experiments and then expand the evaluation on right ventricular strain data from subjects obtained from 3D echocardiographic image sequences, for which different types of strain are available at each point of the right ventricle endocardial surface mesh. Beyond this illustrative application, our methodology has the potential to be generalised to many other population analyses considering heterogeneous high-dimensional descriptors.

Index Terms—Representation learning, uncertainty, information fusion, myocardial strain, cardiac imaging, 3D echocardiography.

I. INTRODUCTION

M. Di Folco, G. Bernardino, P. Clarysse and N. Duchateau are/were with Univ Lyon, Université Claude Bernard Lyon 1, INSA-Lyon, CNRS, Inserm, CREATIS UMR 5220, U1294, F-69621, Lyon, France. E-mail: maxime.difolco@telecom-paris.fr

M. Di Folco was also with the Institute of Machine Learning in Biomedical Imaging, Helmholtz Center Munich, Germany, and is with the LTCI, Telecom Paris, Institut Polytechnique de Paris, France.

G. Bernardino is also with the DTIC, Universitat Pompeu Fabra, Barcelona, Spain.

N. Duchateau is also with the Institut Universitaire de France (IUF).

ACCEPTED FOR PUBLICATION IN IEEE TRANSACTIONS ON MEDICAL IMAGING. DOI: 10.1109/TMI.2026.3698240. © 2026 IEEE. Personal use is permitted. For all other uses, permission must be obtained from IEEE.

DESPITE advanced and relevant processing pipelines, medical image analysis methods suffer many times from two strong limitations: (i) the difficulty of establishing standards for some types of computations, and (ii) the blind acceptance of the input data as it is. Both issues can strongly affect the subsequent computations and therefore the disease analysis. The latter means that a single image is often considered representative of a patient's condition, while uncertainties on acquisition and measurements are well known. In this paper, we explicitly target the former issue, by proposing an original pipeline based on manifold alignment to match latent representations associated to different definitions of a given input descriptor, and exploit the latent space to quantify the discrepancy associated to such definitional differences.

A. Uncertainty modelling on high-dimensional data

Uncertainties can arise at four different stages of medical image analysis with machine learning methods: data collection, data labelling, model selection, and model-based inference [1]. Uncertainty sources are generally distinguished between aleatoric (typically, randomness in the data) and epistemic (namely, systematic) [2]. The latter is directly linked with the design of machine learning models, and therefore receives most of the attention in the literature [3], [4], which also distinguishes between probabilistic and deterministic methods.

In this paper, we consider the different ways a given data descriptor is defined or computed (myocardial strain from a given mesh sequence previously obtained from motion tracking). This notion does not explicitly correspond to aleatoric uncertainties, but also reflects different views of the same patient data. In our cardiac imaging application, the strain patterns correspond to high-dimensional descriptors available at each point of the endocardial surface, and the strain patterns obtained from different definitions correspond to different types of descriptors. In this context, multimodal dimensionality reduction methods are relevant to estimate an intermediate latent space of lower dimensionality in which we can represent these data and the correspondence between the different descriptors [5], which we exploit to consider the definitional divergence between such descriptors.

The use of such intermediate latent space was introduced in geomathematics to estimate possible permeability maps from water flow data, but without explicitly quantifying uncertainty [6]. Explicit estimation of aleatoric uncertainty was obtained

from the variance of high-dimensional descriptors reconstructed from the latent space, for example to localize myocardial infarct from 3D deformation [7] or model the subsurface of the Earth [8]. In the former, uncertainty was modeled globally in the latent space, while the latter adjusted the size of the confidence regions (with probabilistic estimation) locally. In both cases, a single type of high-dimensional input data and (non-deep) manifold learning methods were used, meaning that reconstruction was done a-posteriori. More recently, [9] exploited latent space correspondences between two types of descriptors (images and segmentations) to model uncertainty on such segmentations. The latent space was estimated from an auto-encoder, and latent space correspondences were obtained from the CLIP method [10] originally designed to match image and text embeddings.

Here, we go further by relying on a manifold alignment scheme (Multiple Manifold Learning (MML) [11]–[13]) that not only generalizes the alignment scheme to more than two descriptors, but also provides sample-wise flexibility depending on the link between the input descriptors. Manifold alignment is relevant as some data descriptors may have non-comparable numerical values but encode similar information, which should lead to aligned latent spaces, namely samples with close latent coordinates. Conversely, descriptors that do not carry similar information should lead to samples with more distant latent coordinates. It may be seen as a softer version of point cloud registration [14] and with dimensionality reduction performed at the same time. It pursues objectives comparable to contrastive learning [15], but within a manifold learning perspective, namely providing a latent space into which distances have statistical meaning. This notably allows navigating around existing coordinates as we propose to estimate definitional divergence. Note that our purpose is not to build a consensus across several complementary descriptors, as in many multimodal approaches, most falling under the umbrella of fusion methods [16], recently empowered with the attention mechanisms from transformers [17]. Also, achieving a perfect consensus is not possible, since some of the descriptors may contain partially different information. We therefore hypothesize that the principles of manifold alignment will allow relevant and rather direct sample-wise modelling of definitional divergence in the latent space, conditioned by the degree of matching between the input descriptors.

B. Application: 3D Right Ventricle (RV) strain quantification from echocardiography

The illustrative application we target here concerns the computation of 3D deformation (strain) of the RV from existing mesh sequences. Such computations generally involve a local reference frame based on three orthogonal directions derived from the geometry of the ventricle: radial (from endocardium to epicardium), circumferential (along the circumference), and longitudinal (from apex to base) [18]. This has advantages over using Cartesian coordinates, as this anatomical frame can be related to the directions of the myocardial fibers. However, compared to the Left Ventricle (LV), the RV is particularly challenging because of its asymmetric shape [19]: there is no

consensus on the definition of these local directions for the RV [20], and 3D RV strain quantification remains sensitive to differences in their definition and resulting computations. As a result, for echocardiography, clinical studies focus on 2D RV strain quantified globally or regionally [21], which leads to limited disease characterization. In 3D, coordinates-independent solutions consist in considering area strain (the relative area change of elements of the RV surface) [22], or principal strain obtained by eigendecomposition of the strain tensor [23]. Coordinates-dependent methods can consider global axes aligned to the main RV dimensions [24], or directions that mimic the definition given above [25], [26], but even subtle differences can impact RV local strain [27].

C. Contributions

We propose an original method to consider definitional divergence (potentially involving heterogeneous data types) related to high-dimensional descriptors extracted from medical images. Manifold alignment matches the latent representations associated to the different high-dimensional input descriptors. Definitional divergence is modelled locally in the low-dimensional latent space, and visualized after reconstruction through a high-dimensional parametric map. We demonstrate this approach on experiments with toy and real datasets, using the data from a popular computer vision public dataset, and from a private dataset corresponding to a real cardiac imaging application where the definition of standards for the input descriptors is still a topic of debate [20], [27].

II. METHODOLOGY

Our methodology consists of three consecutive stages, as summarized in Fig.1:

- The joint low-dimensional embedding of high-dimensional descriptors by manifold alignment (Fig.1a and Sec.II-A.1),
- The sampling of new low-dimensional points by estimating a distribution mimicking definitional divergence in the latent space for each subject (Fig.1b and Sec.II-A.2),
- Reconstructing a high-dimensional parametric map from the sampled low-dimensional points (Fig.1c and Sec.II-A.2), referred to as the *discrepancy map*.

The description of this generic approach is then complemented by insights on the 3D RV strain data analyzed in this paper (Sec.II-B). Our code is public¹, and includes a demo corresponding to the experiments on the toy data described in Sec.III-B.

A. Quantification of definitional divergence

1) *Manifold alignment*: Given a population of K subjects, we denote:

- $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_K]^T \in \mathbb{R}^{K \times D}$ the input data corresponding to one high-dimensional descriptor (D being its dimensionality): for example, the images associated to one channel of a set of RGB images as in Sec.III-B, or the RV

¹https://github.com/maxdifolco/definitional_divergence




Fig. 1: Overview of the pipeline proposed in this paper. (a) High-dimensional descriptors (here, 3D patterns corresponding to different definitions of RV strain) are encoded into low-dimensional latent spaces that are locally aligned depending on neighborhood relationships between samples (Sec.II-A.1). (b) Remaining differences in the latent space after alignment are exploited to estimate a statistical distribution that models definitional divergence (Sec.II-A.2). (c) Reconstructing high-dimensional samples from this distribution to visualize such definitional divergence (Sec.II-A.2).

strain patterns associated to a given process to compute strain as in Sec.III-D,

- $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_K]^T \in \mathbb{R}^{K \times d}$ the low-dimensional latent coordinates corresponding to these data, estimated by dimensionality reduction methods, with $d \ll D$.

Given M different descriptors of the same dimensionality D , we denote \mathbf{x}_i^m and \mathbf{z}_i^m the input data and low-dimensional coordinates associated to the m -th descriptor for the i -th subject, with $m \in [1, M]$.

We perform manifold alignment through the MML algorithm² [11]–[13]. It consists in simultaneously estimating a low-dimensional latent space for each descriptor (first term in Eq.1) while controlling, with a hyperparameter μ , the alignment of the latent coordinates (second term in Eq.1). This amounts at minimizing:

$$E(\mathbf{Z}) = \sum_{m=1}^M \sum_{i,j=1}^K \|\mathbf{z}_i^m - \mathbf{z}_j^m\|^2 W_{ij}^m + \mu \sum_{\substack{m,n=1 \\ m \neq n}}^M \|\mathbf{z}_i^m - \mathbf{z}_j^n\|^2 V_{ij}^{mn}, \quad (1)$$

where $\mathbf{W}^m = [W_{ij}^m] \in \mathbb{R}^{K \times K}$ stands for the affinity matrix of the m -th descriptor, and $\mathbf{V}^{mn} = [V_{ij}^{mn}] \in \mathbb{R}^{K \times K}$ encodes the correspondences between the m -th and n -th descriptors, weighted by a factor $\mu > 0$. The elements of both matrices have values between 0 and 1. Note that this formulation requires \mathbf{z}_i^m and \mathbf{z}_j^n to have the same dimensionality.

The first term is similar to the low-dimensional embedding performed by Laplacian eigenmaps on each descriptor [28]: two high-dimensional samples \mathbf{x}_i^m and \mathbf{x}_j^m that are close will lead to an affinity $W_{ij}^m \approx 1$, therefore forcing the low-dimensional coordinates \mathbf{z}_i^m and \mathbf{z}_j^m to be close. Conversely, two high-dimensional samples that substantially differ will lead to an affinity $W_{ij}^m \approx 0$, which will have no influence on bringing \mathbf{z}_i^m and \mathbf{z}_j^m close.

The second term operates in a comparable manner but considering cross-descriptor correspondences through V_{ij}^{mn} , which constrains the distance between the inter-descriptor latent coordinates \mathbf{z}_i^m and \mathbf{z}_j^n , namely the manifold alignment.

In practice, \mathbf{W}^m is defined through a Gaussian kernel, whose width σ is set as the average distance of each point to its k_σ -th neighbor. In addition to μ and k_σ , computations involve one additional hyperparameter, k_M , which stands for the amount of nearest neighbors to sparsify the extra-diagonal matrix \mathbf{V} .

Computational details about the matrices involved in Eq.1 and its solution are given in Appendix A.

2) Modelling definitional divergence: We directly exploit the low-dimensional latent spaces after manifold alignment. Concretely, definitional divergence is modelled sample-wise from the set of M low-dimensional points $\mathcal{Z}^i = \{\mathbf{z}_i^1, \dots, \mathbf{z}_i^M\}$ associated to the i -th subject. If the neighbors of this subject are preserved across the different descriptors, then its low-dimensional coordinates \mathcal{Z}^i should be close. Conversely, if the descriptors differ, such coordinates should be distant, proportionally by how much the descriptors differ.

Our methodology involves a three-stage process. We first estimate a multivariate Gaussian defined by the set of points \mathcal{Z}^i , whose variance and principal axes are obtained in practice by PCA on this set (see illustration in Appendix B). Our experiments involve two or three descriptors, meaning that $|\mathcal{Z}^i| = 2$ or 3, and we therefore consider one or two main directions to define such Gaussian distribution, respectively. Then, we sample N new low-dimensional points according to this distribution ($N = 100$ in all our experiments). This value respectively corresponds to two times the recommended number of samples for estimating Gaussian distributions governed by 2 principal axes with a confidence level of 95% (which corresponds to the case $|\mathcal{Z}^i| = 3$, the recommendation being lower for 1 principal axis, namely $|\mathcal{Z}^i| = 2$) [29]. Finally, we estimate the high-dimensional data associated to each of these new low-dimensional samples. A high-dimensional parametric map reflecting definitional divergence is obtained by dimension-wise standard deviation over the N reconstructed high-dimensional samples (in our case, at each point of the 3D RV mesh).

Any of the M descriptors can be reconstructed, although we mostly make observations on a single descriptor that stands as

²https://gitlab.com/maxDiff/mml_pressure-volume-demo

reference. In our case, we reconstructed the high-dimensional data by regression given that MML does not explicitly have a decoding step. Specifically, we used multiscale kernel ridge regression³ [30] which is robust to the non uniform density of samples in the latent space.

B. Application to 3D RV strain

Our applicative focus is myocardial deformation (3D strain, at each point of a 3D mesh of the RV), and the different definitions or different ways to compute it.

1) Data and pre-processing: We processed RV surface meshes of 100 control subjects obtained from semi-automatic endocardial segmentation by an expert clinician and tracking of 3D transthoracic echocardiographic sequences using commercial software (4D RV Function 2.0, TomTec Imaging Systems GmbH, Germany). More recent versions of this software exist, but this version corresponds to the one used in previous studies led by our clinical collaborator in which the studied subjects served as control cases [22], [26], [31]. The protocol of these studies complied with the declaration of Helsinki and was approved by the local research Ethics committee. All subjects consented and provided written informed consent. The 3D image sequences were acquired from an apical four-chamber view focused on the RV, using a matrix-array X5-1 transducer (Philips Medical Systems). Care was taken to maximize the frame rate (26 ± 11 Hz on this dataset) and to include the entire RV within the images. The tracking software tracked the RV endocardial surface along the cardiac cycle using 3D speckle-tracking [32], and allowed exporting the sequence of 3D meshes for post-processing.

The RV surface meshes were handled as VTK files, and consisted of 822 points and 1587 triangular cells, after cropping out the tricuspid and pulmonary valves. The commercial software uses a mesh model that provides point-to-point mesh correspondences across subjects, attached to the mesh data. We also realigned them across the studied population using generalized Procrustes analysis with a rigid transform.

The RV end-diastolic volume, RV end-systolic volume, and RV ejection fraction were 64.3 ± 25.3 mL, 29.2 ± 13.3 mL, and 55.0 ± 6.3 %. Additional clinical characteristics about these patients can be found in the corresponding studies from our clinical collaborator [22], [26], [31].

2) Local anatomical directions: As for the left ventricle, local directions can be defined at each cell of the RV surface mesh, to be used afterwards to express strain along these directions. The radial direction is defined as the normal to the RV surface at each point. The longitudinal direction can be defined in three manners [27] (Fig.2):

- **Long-axis method:** [22] The long-axis is defined as the segment joining the apex and the basal point equidistant from the valves centers. The circumferential direction is first estimated locally from the cross product between the radial direction and the long-axis. Then, the longitudinal direction is obtained from the cross product between the radial and circumferential directions.

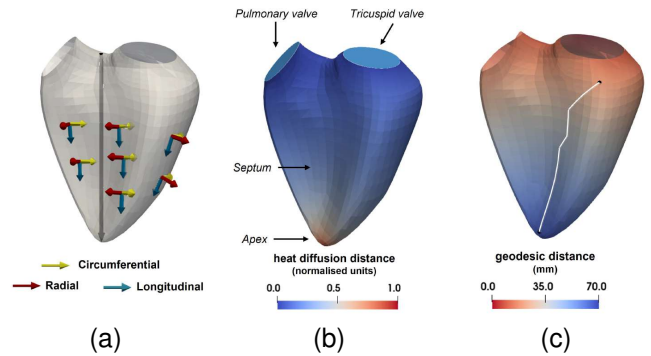


Fig. 2: The three ways to compute circumferential and longitudinal directions we evaluated in this paper. (a) Long-axis computations, (b) Heat diffusion computations (the turquoise disks represent the cold point, while the apex stands as the hot point), (c) Geodesic distance computations (the white line corresponds to the geodesic joining the two purple dots).

- **Heat diffusion method:** [25] The apex and valves are defined as hot ($u = 1$) and cold ($u = 0$) points, respectively. The longitudinal direction is estimated as the gradient of the map u , defined at each point of the mesh by solving the partial differential equation $\nabla \cdot (\nabla u) = 0$. In our implementation, the map u is estimated iteratively by (at each iteration) setting the value at each point as the weighted average of the values at neighboring points in the graph defined by the RV mesh, updating all points, and then restoring the original values 1 and 0 to the apex and the valves.
- **Geodesic distance method:** The longitudinal direction is defined as the gradient of the geodesic distance to the apex. We computed the exact geodesics in the discrete surface [33]. These do not necessarily follow the cell edges, and are therefore less prone to approximations compared to shortest path algorithms such as Dijkstra's.

Finally, the circumferential direction is computed as the cross product between the radial and longitudinal directions.

3) RV strain computations: Once the local directions are computed, the local strain tensor is estimated as the Green Lagrangian strain:

$$\mathbf{E} = \frac{1}{2}(\mathbf{J}^T \cdot \mathbf{J} - \mathbf{I}), \quad (2)$$

where $\mathbf{J} = \nabla v + \mathbf{I}$, with ∇v the displacement gradient at a given point of the RV surface mesh, in Cartesian coordinates, and \mathbf{I} is the identity matrix.

Then, longitudinal and circumferential directional strains are obtained by projecting the strain tensor along these two directions, as:

$$\mathbf{E}_h = \mathbf{h}^T \cdot \mathbf{E} \cdot \mathbf{h}, \quad (3)$$

where \mathbf{h} is the unit vector defining the considered direction. As only endocardial surface meshes were available for export from the commercial software used for segmentation and tracking, radial strain was not considered in this paper.

³<https://github.com/nicolasduchateau/multiscale-kernel-regression>

In our database, all the strain patterns were available at each point of the RV surface and at each instant of the cycle. Nonetheless, we focused the evaluation on end-systolic strain patterns, of higher magnitude. In all figures, results are displayed on end-diastolic meshes, which better render anatomical differences between subjects before deformation.

The algorithms to compute the local anatomical directions and 3D RV strain of a 2D surface embedded in a 3D space were made publicly available⁴ following our publications dedicated to these aspects [27], [34].

4) *Comparison schemes*: We exploited the methodology described in Sec.II-A on RV strain resulting from the following three configurations:

- Longitudinal strain obtained from two different computations of the local anatomical directions,
- Longitudinal strain obtained from three different computations of the local anatomical directions,
- Longitudinal against circumferential strain.

As local directions and therefore strain computations may be influenced by the local RV shape (as we observed earlier [27]), we therefore specifically examined the link between the reconstructed high-dimensional parametric map (referred to as the *discrepancy map*) across the whole RV surface and the shape onto which computation were made. To do so, we used the Partial Least Squares (PLS) algorithm which serves to quickly relate two high-dimensional variables (in our case, the discrepancy map and the RV shape). It provides a low-dimensional space whose main directions maximize the covariance between these two input descriptors, along which we can easily sample representative coordinates across the data distribution and reconstruct their corresponding high-dimensional data, in other words the RV shapes that are most related to the main trends in discrepancy maps.

III. EXPERIMENTS

A. Implementation details

For all experiments, N , the number of points sampled in the latent space to model definitional divergence is set to 100. For the experiment on the COIL-100 dataset (Sec III-B), MML was used with the following parameters: $\mu = 1$, $k_\sigma = 5$, and $k_M = 5$. Concerning the experiment on the RV strain (from Sec III-C to Sec III-G), MML was computed with $\mu = 1$, $k_\sigma = 10$ neighbors and $k_M = 10$. These values are similar to those determined in our previous work [36], which used MML on 3D RV meshes from the same commercial software. This work included a sensitivity analysis of the influence of such hyperparameters on the learnt latent representation for synthetic data (left ventricular pressure and volume curves). In brief, for the RV meshes data, k_σ was set to around 10% of the available samples. At fixed k_M , the optimal μ was chosen as the one minimizing the energy defined in Eq.1. k_M was empirically determined by examining the latent representation.

Computations were fast. On a laptop with Intel Core 4.6 GHz i5-1335U processor and 32 Gb memory, executing the whole pipeline (Matlab code) took less than 10 s for each of

the two objects from the COIL-100 dataset, and for each experiment on RV strain patterns from the 100 control subjects.

B. Toy data: color channels in RGB images

We first evaluated our approach on image samples from the COIL-100 dataset [35] which consists of RGB images of 100 objects, each object being observed under the 72 same viewpoints (a set corresponds to a 360-degree rotation of the object around itself). The three channels of an RGB image can be seen as three different but related ways to represent the observed object, which we may consider as different descriptors of this object. In this experiment, we used our method to apprehend the definitional divergence associated with representing a given object from two of these different channels. The discrepancy map linked to a given channel (e.g. red) corresponds to the error in reconstructing this channel from a latent representation that does not exactly correspond to this channel (e.g. latent coordinates between the red and the green latent coordinates).

We picked two datasets for which the two channels provide complementary information on a large enough portion of the object, such as the soda can (where the red channel is almost uniform, while the green channel is almost present on the brand name) and the toy bear (where the red and green channels are almost complementary except around the nose). We restricted the view angles to prevent redundancies in the images and circular distributions of latent coordinates (from 0° to 90° and 270° to 355° for the toy bear, from 0° to 115° for the soda can, each by steps of 5°). To benefit from more samples and wider distribution of samples in the latent space, we added rotated images of each existing samples (between -5° and $+5^\circ$ by steps of 1°). This led to 264 and 407 samples for the soda can and toy bear, respectively.

Figure 3 displays the original RGB images and analyzed red and green channels at representative viewpoints, and the discrepancy maps estimated for each channel (denoted U_R and U_G), obtained after aligning the red and green channel latent spaces. For the soda can, the two channels do not exactly carry the same information and therefore each channel impacts the reconstruction of the other channel. Indeed, the red channel is rather uniform within the object and therefore represents less the object rotation compared to the green channel. Regarding reconstruction, given the uniformity of the red channel, the highest values in the discrepancy map are mostly concentrated at the borders of the object. Conversely, the green channel mostly contains information near the brand name and the Santa drawing, and the highest values in the discrepancy map are concentrated near the borders of these locations. A comparable interpretation can be given for the toy bear. Both red and green channels carry information about the object rotation, as the orientation of the object within the images (and its pixel intensities) varies across the different samples. However, this concerns the whole object for the red channel, while this only concerns part of the head for the green channel (namely, less accurate information about the whole object rotation). Regarding reconstruction, as the red channel is rather uniform except around the neck, the highest values in the discrepancy

⁴<https://github.com/gbernardino/rvmep>




Fig. 3: Toy experiment involving the red (R) and green (G) channels of images of two objects from the COIL-100 dataset [35], of different view angles and additionally rotated between -5° and $+5^\circ$. (a) Samples of the two objects. The red channel dominates in both objects, being either almost uniform (soda can) or almost complementary of the green channel (toy bear). (b) Estimated discrepancy map associated to definitional divergence, respectively for the red and green channels, for the sample with the largest differences between the red and green latent spaces, after alignment with MML. (c) Similar display for the sample with the smallest differences.

map are mostly around the borders of the object. They are also located around this zone for the green channel, and not around the neck despite some information on this channel, as most of the necklace (except its right extremity) remains unchanged when rotating the view angle.

C. Toy data: RV strain vs. noisy RV strain

We also evaluated our method on another toy dataset, using the RV strain from our medical application. The representation was estimated from the whole population of 100 control subjects. The two descriptors (namely, the high-dimensional observations $\{\mathbf{x}_i^0\}$ and $\{\mathbf{x}_i^1\}$) were the original RV strain (estimated using the *long-axis* method) at each point of the mesh, and the same RV strain but with artificially introduced noise in a specific region of the RV (same region for all subjects). Gaussian noise ($\mathcal{N}(0, 100)$) was added to the original strain values (namely, to the $\{\mathbf{x}_i^1\}$), controlled by a scaling factor α . Figure 4 summarizes our observations. The left part of the figure indicates the zone where noise was added, while the other miniatures depict the discrepancy maps for $\alpha = 0.1, 0.5$, and 1 ($\alpha = 1$ actually means substantial noise compared to the peak strain magnitude around 0.5, see Figs.5 and 6). The values of the discrepancy map increase according to the value of α , mainly in the zone where noise was added. The

non-zero values outside of the noisy zone are both due to the sampling process in the latent space, which induces variations on the whole high-dimensional strain pattern and therefore across the whole RV, and to the regression, which averages neighboring patterns. Despite these limitations, our method is able to capture definitional divergence in a relevant manner for all values of α .

To complement these qualitative observations with quantitative measurements, we generalized this experiment to the 8 different zones into which clinicians commonly divide the RV surface, defined in [37]. We added noise with $\alpha = 1$ in each of these zones, separately, and quantified the average of the discrepancy map per zone for the whole population (Tab.I). The last row corresponds to the average of the discrepancy map obtained between the *long-axis* and *heat diffusion* computations, for comparison purposes. We observe that the estimated values are always more important in the noisy zone compared to the other ones. A global increase is also observed compared to the *long-axis* vs. *heat diffusion* values, mostly because the range of values with $\alpha = 1$ leads to larger strain values.




Fig. 4: Toy experiment involving as high-dimensional observations the original RV strain ($\{\mathbf{x}_i^0\}$) and an artificially noisy version of it ($\{\mathbf{x}_i^1\}$). Noise was introduced on the high-dimensional RV strain in a given zone (see left side of the figure), and we tested different noise intensity levels ($\alpha \in [0, 1]$).

		Average of the discrepancy map per zone							
		#1	#2	#3	#4	#5	#6	#7	#8
Noisy zone	#1	17.56	4.07	2.96	2.32	6.10	3.31	2.53	1.85
	#2	5.62	16.95	4.51	2.68	3.83	3.94	2.80	3.15
	#3	4.36	6.74	17.96	6.50	5.59	6.25	3.50	2.24
	#4	2.79	2.82	6.40	18.04	5.27	4.92	7.03	3.45
	#5	8.35	3.40	3.91	3.86	18.45	6.68	5.92	2.92
	#6	2.60	2.61	3.60	3.05	5.51	17.96	5.04	3.96
	#7	2.10	1.97	2.67	4.70	5.84	6.22	17.57	6.01
	#8	2.13	3.37	2.53	3.25	3.71	5.88	7.03	17.49
Baseline	0.68	0.66	1.05	0.90	1.13	1.35	1.13	0.79	

TABLE I: Toy experiment involving noisy ($\alpha = 1$, in a given zone of the RV) and original RV strain patterns (*long-axis* computation). Values correspond to the average of the discrepancy map, in % (strain units), across the population of 100 control subjects. The baseline corresponds to the discrepancy map estimated between the *long-axis* and *heat diffusion* computations.

RV zones: #1: outflow tract, #2: anterior wall, #3: lateral wall, #4: inferior wall, #5: infundibular septum, #6: membranous septum, #7: inlet septum, #8: trabecular septum.

D. Longitudinal strain

The previous toy experiments demonstrated that our approach is able to identify a relevant zone related to definitional divergence and return higher values in case of higher differences between the input descriptors. In this section, we compare the discrepancy maps associated to the different computations of longitudinal strain described in Sec.II-B: *long-axis*, *heat diffusion* and *geodesic*. We chose the *long-axis* as reference computation and estimated the discrepancy map on this descriptor when also considering one of the two others strain computations as second descriptor, or considering all three together.

Figure 5 displays the strain pattern estimated from the three different types of computation for a representative subject, and the corresponding discrepancy maps.

The strain and discrepancy maps are quite similar in general, except for the zones marked with red (strain) and blue (discrepancy) circles, respectively. Between the valves, the *geodesic* strain is different from the *long-axis* strain, which is correctly captured in the discrepancy map. Similarly, the

heat-diffusion strain pattern is slightly more expanded and pronounced in this zone, as also visible on the discrepancy map between these two descriptors. Comparable observations can be made for the other zones marked with a circle. Finally, the discrepancy map resulting from the analysis of the three strain descriptors together kind of merge the previous observations in a single map.

We then computed the mean discrepancy map across the population for the three different types of strain computation (*long-axis* either vs. *heat diffusion* or *geodesic*, and the three descriptors together), and displayed it on the average mesh for the population, and end-diastole (Fig.6). We discretized the color code to enhance the visualization of distinct zones. We first observe that the displayed maps are quite similar across the three types of strain computation, with the highest values around the valves and near the apex. Local high values are observed under the pulmonary valve in Fig.6b and between the valves in Fig.6c). These also appear when the three descriptors are considered together. This experiment shows that our approach correctly reflects definitional differences between each type of computation when more than two input descriptors are used.

Of note, the high values of the discrepancy map for all of the experiments are concentrated in a zone that is difficult to track in 3D echocardiography (close to the apex and the valves), also meaning that the maps we estimate might be correlated to the local shape of the RV. This is what we specifically evaluate in Sec.III-F.

E. Comparison with the standard deviation between the input descriptors

One may wonder if the manifold alignment and exploitation of the latent space are really necessary, compared to a much simpler approach of directly computing the standard deviation between the input descriptors for each sample (if the different types of such descriptors allow it). This is illustrated in Fig.7, both for the toy data from Sec.III-B and for the longitudinal strain data from Sec.III-D. On the toy data, the discrepancy map was overestimated with this approach, both regarding magnitude and spatial distribution. First, high values are spread almost uniformly across the whole object, as the red and green

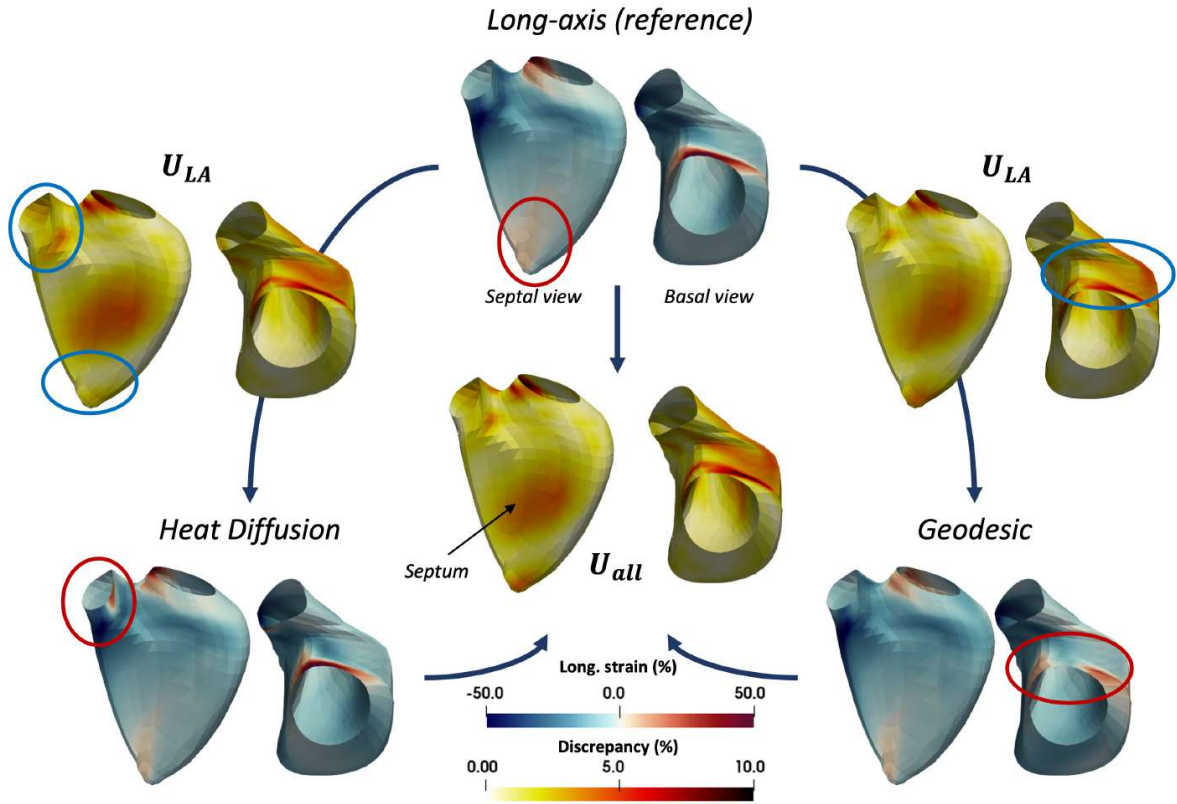


Fig. 5: Strain patterns from a representative subject, and corresponding discrepancy maps. The Long-axis (LA) is used as reference for comparison between two descriptors (with *heat diffusion* and *geodesic* computations) and all three descriptors together (center of the figure). The red and blue circles highlight zones of major differences for the strain and discrepancy maps, respectively.

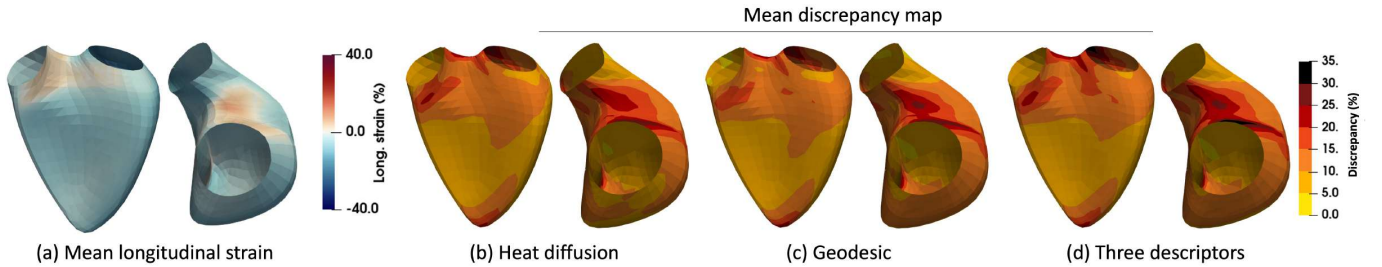


Fig. 6: Average discrepancy map (b-d) across the population for the different strain computation schemes described in Sec.II-B.4. (a) Average longitudinal strain pattern of the *long-axis* computation, used as reference in this experiment.

channels do not have comparable intensities at the same locations (e.g. almost no green color over the can and bear body). However, both channels encode the object rotation (although with slight differences), and the impact of a given channel on the other one should therefore be lower, as previously observed in Fig.3. Besides, in Fig.7, the displayed values have comparable magnitude for the two displayed orientations, while these were actually associated to the largest/smallest discrepancies in Fig.3, due to different grades of manifold alignment around these samples. The strain data is more subtle to assess, but we also observe an overestimation of the discrepancy map both regarding magnitude (e.g. near the valves) and spatial distribution (e.g. near the septal/lateral wall

junction or near the apex).

This simple experiment confirms the relevance of manifold alignment: besides different numerical values, the input descriptors may encode similar information for some samples, which is insufficiently captured by directly computing the standard deviation between the input descriptors.

F. Link with the RV shape

In this section, we examine the link between the discrepancy maps obtained using the three descriptors related to the different strain computations, and the local RV shape, using the experimental setup based on PLS and described in Sec.II-B.4. Shape was considered through the set of 3D coordinates




Fig. 7: Standard deviation between the input descriptors. (a) and (b) Red and green channels, for the soda can and toy bear at the orientations shown in Fig.3b and c. (c) All the three strain descriptors, for the subject depicted in Fig.5.

at each point of the mesh. Figure 8 shows the first three joint modes of variation of shape and definitional divergence, estimated by reconstructing shape and discrepancy maps at -2σ and $+2\sigma$ along each dimension estimated by PLS, with σ the standard deviation of the corresponding dimension.

The first mode of variations mostly encodes the relative position of the valves (tricuspid vs. pulmonary valves). High values in the discrepancy map are observed under the pulmonary valve when the tricuspid valve is more elongated. This mode of variation also encodes the length of the RV, with high values near the apex when the RV is long. The second mode of variation corresponds to global volume and septal curvature (observable in the costal view). An increase of values in the zones marked by the red circles is observed when the curvature increases. Both the second and third modes of variation show larger and more round shapes with high values of the discrepancy map under the valves.

While variations in volumes and lengths are expected with a linear approach like PLS, most shape variations align with regions challenging to track in 3D echocardiography: the apical septum and areas near the valves. Defining anatomical directions in these areas is difficult, resulting in distinct strain patterns between the different computations, leading to high discrepancies.

G. Longitudinal vs. Circumferential strain

In this section, we extend the evaluation to longitudinal and circumferential strain using the *long-axis* computations. In contrast to the previous experiments on the RV, the two descriptors represent different information as visible from the mean strain patterns in Fig.9. Circumferential strain is higher on the free wall and near the border between the septum and the free wall, while longitudinal strain is higher around the valves and the apex, leading to different observations. The U_{Long} pattern is similar to the one observed in Fig.6, indicating that despite considering another descriptor, our method is able to capture the variations of the descriptor individually. Concerning the U_{Circ} pattern, the zone of high




Fig. 8: First three modes of variation of discrepancy maps and RV shapes obtained from PLS. σ stands for the standard deviation along the corresponding dimension.

values corresponds to the area with high strain, which is consistent with the observations made for longitudinal strain.

IV. DISCUSSION

In this paper, we proposed an original methodology to apprehend definitional differences for high-dimensional descriptors extracted from imaging data. Our methodology exploits representation learning and in particular manifold alignment, which puts in correspondence the latent spaces associated to different input descriptors, these descriptors being seen as different versions of the input data. The definitional divergence is considered from the amount of (mis)alignment in the latent space, directly linked to the differences in information carried out by the different descriptors.

We demonstrated the relevance of our method on two toy experiments. First, on some samples from the COIL-100 dataset of RGB images, where the estimated discrepancy map corresponded to differences in the information encoded by some RGB channels. Then, on synthetic data generated from real RV strain patterns with additional noise in a given region (which enables quantitative assessment, and the noisy region can stand as ground truth location), the discrepancy map actually reflected the noise level introduced in such region. Finally, we evaluated the soundness of our method on real RV strain patterns from 100 control subjects, where definitional divergence corresponded to different ways to compute local anatomical coordinates across the RV and therefore myocardial strain. Our approach was able to consider more than two descriptors at once, and the estimated discrepancy maps were consistent with those estimated from pairs of descriptors taken

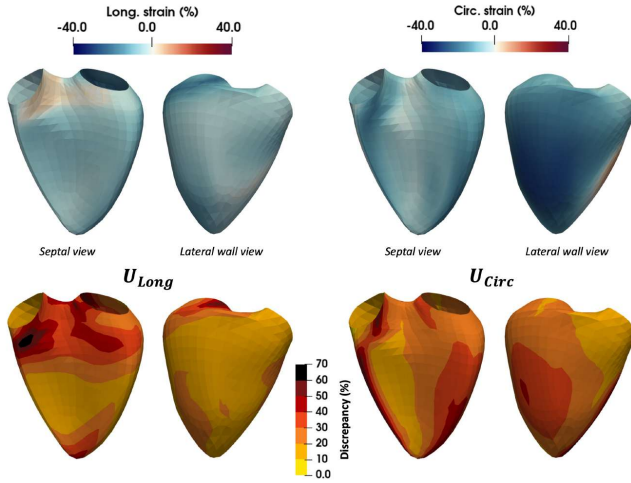


Fig. 9: Average strain (top row) and discrepancy maps (bottom row) for the longitudinal (Long.) and the circumferential (Circ.) strain.

independently, and could be interpreted in light of representative shape variations.

Analysing the mean discrepancy map locally confirms that the highest values were found in regions that are challenging to define anatomical directions, i.e. the apical zone and close to the valve. These regions were identified as the main area of differences when comparing the local differences in terms of anatomical directions and strain patterns [27]. Reaching a consensus on the RV anatomical directions is actually challenging [20]. While the *long-axis* method we took as reference is simpler to compute, the *heat diffusion* method was reported to provide more relevant outputs [27], in comparable computational times. Also note that we focused on different ways to compute myocardial strain from an existing mesh sequence, previously obtained from segmentation and tracking by commercial software. Segmentation and tracking methods, the mesh creation process [38], and the imaging modality and its geometrical characteristics (e.g. 2D or 3D) may also cause additional differences in the quantification of the actual physical strain [39]. Here, we preferred to remain in a more controlled and narrow setting (different definitions of strain applied to 3D echocardiographic data, complemented by experiments on a toy dataset of RGB images) to better focus on the core method to apprehend definitional divergence. For RV echocardiography, there is a clear potential to build higher trust in 3D measurements, which are necessary given that the specific shape of the RV is insufficiently rendered in 2D. This also encourages the analysis of subtle local deformation patterns for disease characterization, beyond the quantification of peak global or wall-specific strain, as highlighted for the LV in 2D [40].

Our methodology can be applied to other types of mesh data, including 3D RV meshes obtained from other cardiac imaging software (provided they can be exported and point-to-point correspondence can be obtained), and to other medical or imaging applications, as illustrated in the toy experiments

in this paper. Application is possible as long as the inputs are of the same size, namely, point-to-point correspondence for meshes and uniform image size for images, and may require adapting the alignment method and metrics used to compare samples (see discussion of these specific aspects further in this section). Extension to pathological cases cannot be done at inference for the representation learning stage, but would instead require learning a new representation that includes such cases. We hope that in the near future such type of computations could be a useful support to clinical interpretations, as already witnessed in other medical imaging fields and a broader scope of uncertainties [41].

We took advantage of PLS to relate the strain discrepancy maps and the 3D RV shape. We preferred this simple approach over non-linear ones, as local shape defects (e.g. folding) may be difficult to highlight on smooth shapes such as the RV from 3D echocardiography [42], [43]. Our analysis indicates that, beyond the anticipated volume and size variations when using a linear approach as PLS, most shape differences and discrepancy maps align with anatomically challenging regions for the RV. In particular, in contrast to the LV, the asymmetrical and crescent shape of the RV introduces distinct challenges. This may question the blind reliance on direction-dependent strain against non-directional one such as area strain, as previously reported [44].

The manifold alignment method we used (MML) suffers from two limitations: the metric to compare samples (here, the Euclidean distance between the data at all mesh points, considered as a column vector), and the reconstruction which is made a-posteriori. We preferred the simpler MML approach to better focus on modelling definitional divergence, which constitutes the core of our methodological contribution. Nonetheless, our methodology is actually adaptable to any type of manifold alignment method, including those based on convolutional networks to better consider the local data structure (which can be seen as using better metrics to compare samples ; our data would require using graph convolutional networks [45]), and those with intrinsic reconstruction such as auto-encoders, which also require being adapted to shape data [46]. Similarly, we preferred to keep the estimation of discrepancy maps from the latent samples simple (localized standard deviation across the reconstructed high-dimensional samples, meaning that each spatial location is considered independently). More advanced methods (e.g. reconstruction within a Bayesian framework) may provide finer estimation of discrepancy maps, in particular with better spatial consistency across the RV mesh, but at the price of higher computational cost and model complexity.

Finally, manifold alignment falls within unsupervised representation learning and is therefore difficult to validate. We therefore designed experiments on toy synthetic and real data for which the locations and magnitude related to definitional divergence can be roughly guessed. Then, we progressively evaluated our methodology on real data with several configurations of increasing complexity (two and three close descriptors, then two correlated-but-different descriptors), thoroughly examining the discrepancy maps with the studied cardiac shapes and our prior knowledge on the data quality and disease.

V. CONCLUSION

We proposed a complete and original pipeline based on manifold alignment to estimate the impact of definitional divergence, namely different ways to compute a given descriptor, or complementary descriptors to represent a single patient's condition. Our approach enables relevant assessment locally, and is easily extendable to more elaborated types of manifold alignment based on graph neural networks and auto-encoders, and many other population analyses relying on heterogeneous high-dimensional descriptors.

APPENDIX A: COMPUTATIONAL DETAILS OF THE MANIFOLD ALIGNMENT

We considered standard formulations for the matrices \mathbf{W}^m and \mathbf{V}^{mn} , namely:

$$W_{ij}^m = \exp \frac{-\|\mathbf{x}_i^m - \mathbf{x}_j^m\|^2}{\sigma^2}, \quad (4)$$

where σ is the width of the kernel defining the affinity matrix, and

$$V_{ij}^{mn} = \frac{\langle \mathbf{w}_i^m, \mathbf{w}_j^n \rangle}{\|\mathbf{w}_i^m\| \|\mathbf{w}_j^n\|} \in [0, 1], \quad (5)$$

where \mathbf{w}_i^m is the i -th row of the affinity matrix \mathbf{W}^m . Variants and discussion around such formulations (including alternative measures of similarity between samples, in particular for the inter-modality correspondences) can be found in [13], [36]. This formulation is somehow comparable to angular-based similarity matrices used for inter-modality correspondences, as proposed in multimodal clustering approaches [47].

Equation 1 can be reformulated using a block matrix $\mathbb{W} \in \mathbb{R}^{MK \times MK}$, whose diagonal and extra-diagonal blocks are the matrices \mathbf{W}^m and \mathbf{V}^{mn} , respectively (see the demonstration in the Supplementary Material of [36]). This matrix formulation leads to solving:

$$\left. \begin{aligned} E(\mathbf{Z}) &= \text{tr}(\mathbf{Z}^T \mathbb{L} \mathbf{Z}) \\ \text{s.t. } \mathbf{Z}^T \mathbf{D}_{\mathbb{W}} \mathbf{Z} &= \mathbf{I}, \end{aligned} \right\} \quad (6)$$

where $\mathbb{L} = \mathbf{D}_{\mathbb{W}} - \mathbb{W}$ and $\mathbf{D}_{\mathbb{W}}$ is a diagonal matrix such that $\mathbf{D}_{\mathbb{W}ii} = \sum_j \mathbb{W}_{ij}$. Eq.6 amounts at solving the generalized eigenvalue problem $\mathbb{L}\mathbf{f} = \lambda \mathbf{D}_{\mathbb{W}}\mathbf{f}$, where λ and \mathbf{f} stand for the eigenvalues and eigenvectors, respectively. In practice, we solve $\mathbb{P}\mathbf{f} = (1 - \lambda)\mathbf{f}$, where $\mathbb{P} = \mathbf{D}_{\mathbb{W}}^{-1/2} \mathbb{W} \mathbf{D}_{\mathbb{W}}^{-1/2}$ is symmetric, which corresponds to working with the normalized graph Laplacian.

The coordinates \mathbf{Z} correspond to the first eigenvectors associated to the first eigenvalues sorted by ascending order after removing the trivial case associated to the eigenvalue zero. As $\mathbb{W} \in \mathbb{R}^{MK \times MK}$, we obtain eigenvectors whose rows $[K(m-1)+1, Km]$ correspond to the low-dimensional coordinates for the m -th descriptor (the method therefore provides one latent space for each descriptor).

APPENDIX B: MODELLING DEFINITIONAL DIVERGENCE IN THE LATENT SPACE

PCA estimates orthogonal axes corresponding to the principal directions of variance in the data. Figure 10 illustrates

its use to estimate a multivariate Gaussian defined by the set of points \mathcal{Z}^i , with $|\mathcal{Z}^i| = 2$ (Fig.10a) or 3 (Fig.10b), and to sample $N = 100$ new points from this distribution.

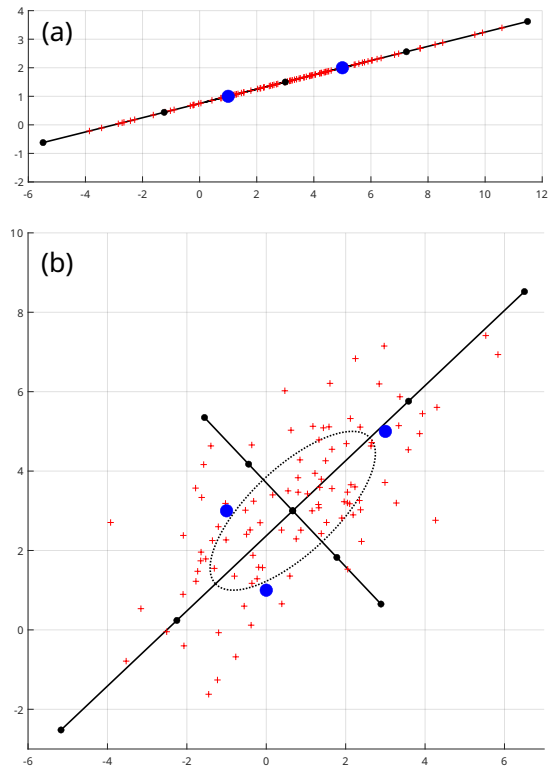


Fig. 10: Estimation of the principal axes (black lines / the black dots correspond to -3 to $+3$ standard deviations along these axes) of a Gaussian distribution from 2 (a) and 3 (b) samples (blue dots), and $N = 100$ random samples generated from this distribution (red crosses). In this illustrative experiment, the horizontal and vertical axes correspond to arbitrary dimensions. In our application, these could be seen as a projection of the low-dimensional samples into the hyperplane defined by the blue dots.

ACKNOWLEDGMENTS

The authors acknowledge the support from the French ANR (LABEX PRIMES of Univ. Lyon [ANR-11-LABX-0063] and the JCJC project “MIC-MAC” [ANR-19-CE45-0005]). Gabriel Bernardino was also partially supported by the FSE+ and “Agencia Estatal de investigación” MICIU/AEI/ 10.13039/501100011033 (RYC2022-035960-I, PID2023-149959OA-I00). They are also grateful to P. Mocerì (CHU Nice, France) for providing the imaging data related to the studied population, and to T. Dargent (LMD, IPSL Paris, France) for the initial computations on local coordinates and 3D RV strain.

REFERENCES

- [1] E. Begoli, T. Bhattacharya, and D. Kusnezov, “The need for uncertainty quantification in machine-assisted medical decision making,” *Nat Mach Intell*, vol. 1, pp. 20–3, 2019.
- [2] E. Hüllermeier and W. Waegeman, “Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods,” *Machine Learning*, vol. 110, pp. 457–506, 2021.
- [3] A. Kendall and Y. Gal, “What uncertainties do we need in bayesian deep learning for computer vision?” *Proc. NeurIPS*, vol. 30, 2017.

- [4] S. Seoni, V. Jahmunah, M. Salvi, P. Barua, F. Molinari, and U. Acharya, "Application of uncertainty quantification to artificial intelligence in healthcare: A review of last decade (2013–2023)," *Comput Biol Med*, vol. 165, p. 107441, 2023.
- [5] K. Kontolati, D. Loukrezis, D. Giovanis, L. Vandanapu, and M. Shields, "A survey of unsupervised learning methods for high-dimensional uncertainty quantification in black-box-type problems," *J Comput Phys*, vol. 464, p. 111313, 2022.
- [6] J. Caers, K. Park, and C. Scheidt, "Modeling uncertainty of complex earth systems in metric space," *Handbook of Geomathematics, Springer Berlin Heidelberg*, p. 865–89, 2010.
- [7] N. Duchateau, M. De Craene, P. Allain, E. Saloux, and M. Sermesant, "Infarct localization from myocardial deformation: Prediction and uncertainty quantification by regression from a low-dimensional space," *IEEE Trans Med Imaging*, vol. 35, pp. 2340–52, 2016.
- [8] R. Blanc and G. Székely, "Confidence regions for statistical model based shape prediction from sparse observations," *IEEE Trans Med Imaging*, vol. 31, pp. 1300–10, 2015.
- [9] T. Judge, O. Bernard, M. Porumb, A. Chartsias, A. Beqiri, and P. Jodoin, "CRISP - Reliable uncertainty estimation for medical image segmentation," *Proc. MICCAI, LNCS*, vol. 13438, p. 492–502, 2022.
- [10] A. Radford, J. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal *et al.*, "Learning transferable visual models from natural language supervision," *Proc. ICML*, vol. 139, pp. 8748–63, 2021.
- [11] J. Valencia-Aguirre, A. Álvarez Meza, G. Daza-Santacoloma, C. Acosta-Medina, and C. Castellanos-Domínguez, "Multiple manifold learning by nonlinear dimensionality reduction," *Proc. CIARP, LNCS*, vol. 7042, pp. 206–13, 2011.
- [12] C. Lee, A. Elgammal, and M. Torki, "Learning representations from multiple manifolds," *Pattern Recognit*, vol. 50, pp. 74–87, 2016.
- [13] J. Clough, D. Balfour, G. Cruz, P. Marsden, C. Prieto, A. Reader *et al.*, "Weighted manifold alignment using wave kernel signatures for aligning medical image datasets," *IEEE Trans Pattern Anal Mach Intell*, vol. 42, pp. 988–97, 2019.
- [14] S. Monji-Azad, J. Hesser, and N. Löw, "A review of non-rigid transformations and learning-based 3D point cloud registration methods," *ISPRS J Photogramm Remote Sens*, vol. 196, pp. 85–72, 2023.
- [15] P. Le-Khac, G. Healy, and A. Smeaton, "Contrastive representation learning: A framework and review," *IEEE Access*, vol. 8, pp. 193907–34, 2020.
- [16] T. Baltrušaitis, C. Ahuja, and L. Morency, "Multimodal machine learning: A survey and taxonomy," *IEEE Trans Pattern Anal*, vol. 41, pp. 423–43, 2018.
- [17] P. Xu, X. Zhu, and D. Clifton, "Multimodal learning with transformers: A survey," *IEEE Trans Pattern Anal*, vol. 45, pp. 12113–132, 2023.
- [18] J. D'hooge, A. Heimdal, F. Jamal, T. Kukulski, B. Bijnens, F. Rademakers *et al.*, "Regional strain and strain rate measurements by cardiac ultrasound: principles, implementation and limitations," *Eur J Echocardiogr*, vol. 1, pp. 154–70, 2000.
- [19] J. Sanz, D. Damián Sánchez-Quintana, E. Bossone, H. Bogaard, and R. Naeije, "Anatomy, function, and dysfunction of the right ventricle: JACC state-of-the-art review," *J Am Coll Cardiol*, vol. 73, pp. 1463–82, 2019.
- [20] N. Duchateau, P. Mocerri, and M. Sermesant, "Direction-dependent decomposition of 3D right ventricular motion: beware of approximations," *J Am Soc Echocardiogr*, vol. 34, pp. 201–3, 2021.
- [21] L. Badano, T. Koliass, D. Muraru, T. Abraham, G. Aurigemma, T. Edvardsen *et al.*, "Standardization of left atrial, right ventricular, and right atrial deformation imaging using two-dimensional speckle tracking echocardiography: a consensus document of the EACVI/ASE/Industry Task Force to standardize deformation imaging," *Eur Heart J Cardiovasc Imaging*, vol. 1, pp. 591–600, 2018.
- [22] P. Mocerri, N. Duchateau, D. Baudouy, E. Schouver, S. Leroy, F. Squara *et al.*, "Three-dimensional right-ventricular regional deformation and survival in pulmonary hypertension," *Eur Heart J Cardiovasc Imaging*, vol. 19, pp. 450–8, 2018.
- [23] A. Satriano, P. Pournazari, N. Hirani, D. Heltersen, M. Thakrar, J. Weatherald *et al.*, "Characterization of right ventricular deformation in pulmonary arterial hypertension using three-dimensional principal strain analysis," *J Am Soc Echocardiogr*, vol. 32, pp. 385–93, 2019.
- [24] M. Tokodi, L. Staub, A. Budai, B. Lakatos, F. Csákvári, M. Suhai *et al.*, "Partitioning the right ventricle into 15 segments and decomposing its motion using 3D echocardiography-based models: The updated ReVISION method," *Front Cardiovasc Med*, vol. 8, p. 622118, 2021.
- [25] R. Doste, D. Soto-Iglesias, G. Bernardino, A. Alcaine, R. Sebastian, S. Giffard-Roisin *et al.*, "A rule-based method to model myocardial fiber orientation in cardiac biventricular geometries with outflow tracts," *Int J Numer Method Biomed Eng*, vol. 35, p. e3185, 2019.
- [26] P. Mocerri, N. Duchateau, S. Gillon, L. Jaunay, D. Baudouy, F. Squara *et al.*, "Three-dimensional right ventricular shape and strain in congenital heart disease patients with right ventricular chronic volume loading," *Eur Heart J Cardiovasc Imaging*, vol. 22, pp. 1174–81, 2021.
- [27] M. Di Folco, T. Dargent, G. Bernardino, P. Clarysse, and N. Duchateau, "Which anatomical directions to quantify local right ventricular strain in 3D echocardiography?" *Proc. FIMH, LNCS*, vol. 13958, pp. 607–15, 2023.
- [28] M. Belkin and P. Niyogi, "Laplacian eigenmaps for dimensionality reduction and data representation," *Neural Comput*, vol. 15, pp. 1373–96, 2003.
- [29] J. Psutka and J. Psutka, "Sample size for maximum-likelihood estimates of gaussian model depending on dimensionality of pattern space," *Pattern Recognit*, vol. 91, pp. 25–33, 2019.
- [30] N. Duchateau, M. De Craene, M. Sitges, and V. Caselles, "Adaptation of multiscale function extension to inexact matching: application to the mapping of individuals to a learnt manifold," *Proc. GSI, LNCS*, vol. 8085, pp. 578–86, 2013.
- [31] P. Mocerri, N. Duchateau, B. Sartre, D. Baudouy, F. Squara, M. Sermesant *et al.*, "Value of 3d right ventricular function over 2d assessment in acute pulmonary embolism," *Echocardiography*, vol. 38, pp. 1694–701, 2021.
- [32] D. Muraru, V. Spadotto, A. Cecchetto, G. Romeo, P. Aruta, D. Ermacora *et al.*, "New speckle-tracking algorithm for right ventricular volume analysis from three-dimensional echocardiographic data sets: validation with cardiac magnetic resonance and comparison with the previous analysis tool," *Eur Heart J Cardiovasc Imaging*, vol. 17, pp. 1279–89, 2016.
- [33] J. Mitchell, D. Mount, and C. Papadimitriou, "The discrete geodesic problem," *SIAM J Comput*, vol. 16, pp. 647–68, 1987.
- [34] G. Bernardino, T. Dargent, O. Camara, and N. Duchateau, "Stranger things: discrete differential geometry for transporting right ventricular deformation across surface meshes," *Proc. FIMH, LNCS*, vol. 13958, pp. 338–46, 2023.
- [35] S. A. Nene, S. K. Nayar, and H. Murase, "Columbia object image library (coil-20)," 1996.
- [36] M. Di Folco, P. Mocerri, P. Clarysse, and N. Duchateau, "Characterizing interactions between cardiac shape and deformation by non-linear manifold learning," *Med Image Anal*, vol. 75, p. 102278, 2022.
- [37] F. Haddad, S. Hunt, D. Rosenthal, and D. Murphy, "Right ventricular function in cardiovascular disease, Part I: Anatomy, physiology, aging, and functional assessment of the right ventricle," *Circulation*, vol. 117, pp. 1436–48, 2008.
- [38] A. Kwan, E. Chang, I. Jain, J. Theurer, X. Tang, N. Francisco *et al.*, "Deep learning-derived myocardial strain," *JACC Cardiovasc Imaging*, vol. 17, pp. 715–25, 2024.
- [39] D. Muraru, U. Cucchini, S. Mihăilă, M. Miglioranza, P. Aruta, G. Cavalli, others *et al.*, "Left ventricular myocardial strain by three-dimensional speckle-tracking echocardiography in healthy subjects: reference values and analysis of their physiologic and technical determinants," *J Am Soc Echocardiogr*, vol. 27, pp. 858–71, 2014.
- [40] N. Duchateau, F. Loncaric, M. Cikes, A. Doltra, M. Sitges, and B. Bijnens, "Variability in the assessment of myocardial strain patterns: implications for adequate interpretation," *Ultrasound Med Biol*, vol. 46, pp. 244–54, 2020.
- [41] B. Kompa, J. Snoek, and A. Beam, "Second opinion needed: communicating uncertainty in medical machine learning," *NPJ Digit Med*, vol. 4, p. 4, 2021.
- [42] S. Jia, N. Duchateau, P. Mocerri, M. Sermesant, and X. Pennec, "Parallel transport of surface deformations from pole ladder to symmetrical extension," *ShapeMI-MICCAI, LNIP*, vol. 11167, pp. 116–24, 2018.
- [43] N. Guigui and X. Pennec, "Parallel transport, a central tool in geometric statistics for computational anatomy: Application to cardiac motion modeling," *Handbook of Statistics*, vol. 46, pp. 28–326, 2022.
- [44] B. Smith, G. Dobson, D. Dawson, A. Charalampopoulos, J. Grapsa, and P. Nihoyannopoulos, "Three-dimensional speckle tracking of the right ventricle: toward optimal quantification of right ventricular dysfunction in pulmonary hypertension," *J Am Coll Cardiol*, vol. 64, pp. 41–51, 2014.
- [45] M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: going beyond Euclidean data," *IEEE Signal Processing Magazine*, vol. 34, pp. 18–42, 2017.
- [46] A. Ranjan, T. Bolkart, S. Sanyal, and M. Black, "Generating 3D faces using convolutional mesh autoencoders," *Proc. ECCV*, pp. 704–20, 2018.

- [47] Y. Wang, X. Lin, L. Wu, W. Zhang, and Q. Zhang, "Exploiting correlation consensus: Towards subspace clustering for multi-modal data," *Proc. ACM MM*, pp. 981–9, 2014.