# Analysis of Nonstandardized Stress Echocardiography Sequences Using Multiview Dimensionality Reduction

*Email address:* mariana.nogueira@upf.edu (Mariana Nogueira)

Accepted for publication in Medical Image Analysis, doi: 10.1016/j.media.2019.101594

## Analysis of Nonstandardized Stress Echocardiography Sequences Using Multiview Dimensionality Reduction

Mariana Nogueira<sup>a,b,\*</sup>, Mathieu De Craene<sup>a</sup>, Sergio Sanchez-Martinez<sup>c</sup>, Devyani Chowdhury<sup>d</sup>, Bart Bijnens<sup>b,e</sup>, Gemma Piella<sup>c</sup>

<sup>a</sup>Medisys, Philips Research Paris, France <sup>b</sup>PhySense, ETIC, Universitat Pompeu Fabra, Barcelona, Spain <sup>c</sup>SIMBIOsys, ETIC, Universitat Pompeu Fabra, Barcelona, Spain <sup>d</sup>Cardiology Care for Children, Lancaster, Pennsylvania, USA <sup>e</sup>ICREA, Barcelona, Spain

## Abstract

Alternative stress echocardiography protocols such as handgrip exercise are potentially more favorable towards large-scale screening scenarios than those currently adopted in clinical practice. However, these are still underexplored because the maximal exercise levels are not easily quantified and regulated, requiring the analysis of the complete data sequences (thousands of images), which represents a challenging task for the clinician. We propose a framework for the analysis of these complex datasets, and illustrate it on a handgrip exercise dataset including complete acquisitions of 10 healthy controls and 5 ANT1 mutation patients (1377 cardiac cycles). The framework is based on an unsupervised formulation of multiple kernel learning, which is used to integrate information coming from myocardial velocity traces and heart rate to obtain a lower-dimensional representation of the data. Such simplified representation is then explored to discriminate groups of response and understand the underlying pathophysiological mechanisms. The analysis pipeline involves the reconstruction of population-specific signatures using multiscale kernel regression, and the clustering of subjects based on the trajectories defined by their projected sequences. The results confirm that the proposed framework is able to detect distinctive clusters of response and to provide insight regarding the underlying pathophysiology.

*Keywords:* Multiview dimensionality reduction, multiple kernel learning, stress echocardiography, pattern analysis

## 1. Introduction

#### 1.1. Clinical Context and Motivation

Stress echocardiography can unveil early-stage cardiovascular-pathology signatures that are not expressed at baseline condition, thus being a valuable tool for screening purposes. Current stress echocardiography protocols, based on exercise or pharmacological stress (Voigt et al., 2003; Davidavicius et al., 2003), are standardized, meaning that the control of the stress levels over the test is very rigorous (based on dose, heart rate, time, etc.), allowing the evaluation of response to stress to be performed based on the comparison of measurements collected at a few discrete timepoints (corresponding to very precise stress levels). However, this standardization comes at the cost of cumbersome protocols, being time-consuming as well as requiring highly-trained staff and specialized equipment. All this translates into high costs, which limit the application of current protocols to a fairly lesser extent than desired, and thus making them unsuited for large-scale screening purposes. Moreover, by getting data at pre-determined intervals and timings, one might be missing pertinent information, as the disregarded dynamic data potentially contain additional valuable information concerning the

<sup>\*</sup>Corresponding author at: DTIC, Universitat Pompeu Fabra (office 55.121), Roc Boronat 138, 08018 Barcelona, Spain.

 $Email \ address: \verb"mariana.nogueira@upf.edu" (Mariana Nogueira)$ 

patient's physiological state.

Other less standard forms of stress, such as the cold pressor test (Velasco et al., 1997) and handgrip exercise (Strauss et al., 2013; Kivowitz et al., 1971; Helfant et al., 1971), were already reported to trigger cardiovascular responses that could unmask differential responses to stress by healthy and pathological patients. These protocols are cheap and practical, come with low risks to the patient, and involve little patient motion, making imaging an easier task. As such, they hold great potential for screening, overcoming the main limitations of current protocols. Besides their potential for screening, they also represent an alternative for patients that are physically unable to undergo a classical exercise test. However, there is one main drawback: the level of exercise is hard to quantify and control, and the timings and magnitudes of events are unpredictable; in other words, they are nonstandardizable. In practice, this implies continuously analyzing the complete acquisition, and focusing on trends/patterns of response rather than on a discrete set of values. The analysis, of these long, dynamic, heterogeneous sequences, which also implies integration of multiple features, is not trivial, and clinicians lack tools to assist them in this task. On the other hand, this type of analysis may be advantageous: by allowing the identification of variations in exercise performance throughout the dynamic range (versus at discrete points in time), it may be more informative of the patients physiological state, and thus have a higher predictive value of adverse outcomes.

Currently, machine learning is being established as one of the preferred tools for the analysis of patterns in functional and high-dimensional data, and has become remarkably popular within the biomedical field. It has already been applied to the study of cardiac response to stress, based on multiple heterogeneous descriptors, such as the velocity profiles of different myocardial segments and timings of key events in the cardiac cycle (Sanchez-Martinez et al., 2017, 2018). However, to the best of our knowledge, it has not yet been used to explore nonstandardized continuous echocardiographic recordings. In this paper, we propose an analysis framework that explicitly addresses the practical challenges this kind of sequences pose, and illustrate its potential in a specific group of cardiac patients.

#### 1.2. Technical Context

In biomedical research, there is an emergent need for machine learning algorithms able to learn from multiple concurrent data sources (e.g. imaging, signal, patient metadata). This type of learning is commonly referred to as multiview learning (Xu et al., 2013). In the cardiac domain, both supervised and unsupervised multiview learning algorithms have been recently applied in the analysis of cardiac motion patterns for numerous applications, e.g. in the identification of dilated cardiomyopathy (Puyol-Antón et al., 2019), in cardiac resynchronisation therapy response prediction (Peressutti et al., 2017) or in the study of heart failure with preserved ejection fraction (HF-PEF) (Sanchez-Martinez et al., 2017, 2018).

In this work, we integrate information coming from multiple heterogeneous features (i.e., heart rate and velocity traces from echocardiographic images) to evaluate patterns of response to stress. Since nonstandardized sequences typically last 60-120 cardiac cycles (equivalent to thousands of images), we propose unsupervised multiview dimensionality reduction to obtain a compact representation of the patterns of response over time. This low-dimensional representation can be used to obtain the principal modes of variation – which describe how the features change – and the temporal trajectories – which encode the timings and intensity of such changes.

Unsupervised multiview dimensionality reduction is an active field of research, including canonical correlation analysis (Hotelling, 1936), partial least squares (Wold, 1985), multiple kernel learning (MKL) (Lin et al., 2011) or multi-modal autoencoders (Li et al., 2018) as some of the most popular algorithms. Our choice for MKL was based on (1) its ability to address inherent nonlinearities of the data and any number of desired input features, without strong assumptions on their correlations, and (2) its good performance in similar applications, while providing a fairly simpler, very flexible, potentially more intuitive/interpretable framework than other types of machine learning.

Once a low-dimensional embedding is estimated, the main modes of variation in the data can be reconstructed using multiscale kernel regression (MKR) (Bermanis et al., 2013; Duchateau et al., 2013). A combined analysis using MKL and MKR was successfully explored before by Sanchez-Martinez et al. (2017, 2018) to characterize functional responses to semi-supine bicycle exercise of controls and patients with heart failure with preserved ejection fraction (HFPEF), based on leftventricular velocity patterns. This work dealt, however, with only two-timepoint (rest/stress) information for each patient, acquired during a standardized exercise stress test.

We propose a technical framework that extends this analysis to the challenging context of nonstandardized stress echo datasets.

#### 1.3. Proposed Approach

Our framework uses MKL to project heterogeneous data collected at each cardiac cycle throughout the stress test onto a low-dimensional space where the main variations in the data are encoded. In this space, the response to stress of each subject can be seen as a trajectory and, based on the similarity among trajectories, subjects can be grouped in clusters that reflect differential patterns of response. The physiological interpretation of the results is decoded through MKR, which allows reconstructing the input signals along any path over the low-dimensional output space.

A preliminary version of the framework was previously proposed (Nogueira et al., 2017). The present paper extends the work in several aspects: we test the framework against a real dataset including healthy and pathological cases, whereas previously the cases had been generated synthetically; we explore other physiological features, using velocity traces at the basal septum of the left ventricle instead of the global longitudinal strain; we reformulate the clustering analysis in the trajectory space by exploring a more sophisticated way of computing distances among trajectories, involving Dynamic Time Warping (DTW) (Bemdt and Clifford, 1994). In addition, we enrich the analysis by exploring and interpreting the spatial configurations of the distributions of the control and diseased population samples in the output space.

## 1.4. ANT1 mutation

To illustrate the framework, we apply it to the discriminative analysis between the dynamics of response to stress in patients with Adenine Nucleotide Translocator-1 (ANT1) deficiency (due to a mutation in an encoding gene) and controls, during handgrip exercise challenges. In patients with ANT1 mutation there is a lack of adenine nucleotide transferase, which converts ADP to ATP. The decreased availability of ATP to the muscles causes lactic acidosis. These patients present with shortness of breath with exercise at a very young age. Within the scope of this paper, they can be considered as extreme cases of HFPEF.

## 2. Methods

A diagram illustrating the main blocks of the framework is depicted in Figure 1. The first block corresponds to the automated processing and extraction of features from the sequence data (Section 2.2). The second block refers to the application of MKL to obtain a low-dimensional representation of the data (Section 2.3). Finally, the third block corresponds to the analysis of this low-dimensional representation, focusing on the discrimination between groups of response and the understanding of the underlying pathophysiological mechanisms (Section 2.4).

## 2.1. Data

This study includes 15 subjects, 10 controls (average age  $24 \pm 14$  years) and 5 ANT1 mutation patients (average age  $21 \pm 7$  years). The echocardiographic acquisitions were performed using a Vivid Q system (GE Healthcare). For each subject, a Doppler myocardial velocity imaging (DMI) sequence of the apical 4-chamber view was acquired (average sampling rate  $115 \pm 43$  Hz) during handgrip exercise. All sequences comprise the start of exercise, a phase of sustained exercise and recovery (average heart rate  $92 \pm 18$  bpm for controls;  $118\pm23$  bpm for ANT1 patients). The durations of each phase vary across subjects. When the 15 subjects are considered, our dataset amounts to a total of 1377 cardiac cycles (average sequence length  $92 \pm 26$  cardiac cycles).

#### 2.2. Feature Extraction

In our dataset, we have an average of about 70 DMI frames per cardiac cycle. As such, 1377 cycles contain a considerably large amount of data, calling for integration and simplification. The first simplification comes with feature extraction, i.e., collecting relevant descriptors of cardiac function throughout the acquisition, while ensuring their robustness to noise and artifacts in the data (e.g. due to breathing or transducer motion). Features should be easy to obtain in clinical practice and, ideally, in an automated manner (manually processing these many cardiac cycles would be impractical).

We selected the left-ventricular basal-septum velocity profile and heart rate (HR) as the features



Figure 1: Main stages of the proposed framework: *i*) automated cycle-wise feature extraction; *ii*) multiview dimensionality reduction to project stress echo sequences onto a low-dimensional space; *iii*) physiological interpretation of the output-space sample distributions and cluster analysis in the trajectory space. DMI=Doppler Myocardial Velocity Imaging. HR=Heart Rate. MKL=Multiple Kernel Learning. DTW=Dynamic Time Warping. MKR=Multiscale Kernel Regression.

of interest to monitor during the stress protocol. These were automatically extracted with the aid of the ECG as temporal reference: the full-length velocity traces were extracted from the DMI sequences using a commercial software (EchoPAC, v.113, GE Healthcare), by manually placing a region of interest (default dimensions) at the basal septal region (see Figure 2a; the actual trace is computed and exported through the software). The cycle-wise traces were obtained by slicing the fulllength profiles at the R-peak positions of the simultaneously acquired ECG and the HR was obtained from the timings of the R peaks (whole process illustrated in stage I of Figure 1). Examples of a full-length velocity trace and sliced cycles (timenormalized for cycle duration as explained in 2.3) are featured in Figure 2b.

Finally, we fed a set of 1377 multiview samples (corresponding to all cardiac cycles of all 15 subjects) to the MKL algorithm, describing each cardiac cycle of each patient by a velocity curve and a HR value (see stage *II* of Figure 1).

#### 2.3. Computation of the low-dimensional space using MKL

Given a high-dimensional dataset with N samples  $X = \{x_i \in \mathbb{R}^d\}_{i=1}^N$ , graph embedding aims at finding a low-dimensional projection  $Y = \{y_i \in \mathbb{R}^k\}_{i=1}^N, k < d$ , that preserves the main topology and variability of the data while removing noisy contributions. To achieve this, a similarity matrix W defined over the pairs of input samples is used to weight the optimization problem which, under appropriate constraints, can be generically expressed as

$$\min_{Y} \sum_{ij} \|y_i - y_j\|^2 W_{ij} \quad . \tag{1}$$

In this way, to minimize the product  $||y_i - y_j||^2 W_{ij}$ , close samples in the input space (high  $W_{ij}$ ) are enforced to remain close in the output space (small  $||y_i - y_j||$ ), while distant samples have little or no influence on each other's optimal projection.



Figure 2: Extraction of velocity sequence data. (a) Example of a frame from a DTI sequence from an ANT1 patient. The yellow circle over the basal septum is the region of interest used to monitor the velocity over the whole sequence. (b) Top: example of a full-length velocity trace. Bottom: isolated rest (left), peak-stress (middle) and recovery (right) cycles, extracted from the corresponding annotated regions in the top plot.

Based on this graph embedding framework, Lin et al. (Lin et al., 2011) generalized the concept of MKL, originally formulated within the support vector machine framework (Bach et al., 2004; Hearst et al., 1998) for classification/regression, to (supervised and unsupervised) dimensionality reduction. By combining multiple kernels, each one based on a specific data descriptor, MKL fuses heterogeneous information and provides the contribution of each feature to the low-dimensional output representation. The unsupervised formulation, adopted in this work, can be summarized as follows.

Let the input dataset, composed of N samples with M descriptors each, be defined as  $X = \{x_i\}_{i=1}^N, x_i = \{x_i^m \in \mathbb{R}^{d_m}\}_{m=1}^M$  where  $x_i^m$  represents the descriptor m associated with sample i and of dimensionality  $d_m$ . The projection of a sample is parametrized by a projection matrix  $A \in \mathbb{R}^{N \times k}$  (where k refers to the selected dimensionality of the output space,  $k \in [1, N - 1]$ ) and a vector  $\beta \in \mathbb{R}^M$  that determines the normalized weight of each feature in the mapping. A unified mapping based on heterogeneous descriptors is made possible as A and  $\beta$  operate on kernelized data rather than on their raw content. For each feature, a kernel matrix  $K_m$  is defined, encoding the similarities over the pairs of samples, based on kernel functions  $k_m$ , i.e.,

$$K_m \in \mathbb{R}^{N \times N}$$
 with  $K_m(i,j) = k_m(x_i^m, x_j^m)$ .  
(2)

In this work,  $k_m$  is a Gaussian kernel (with Euclidean distance) whose bandwidth  $\sigma_m$  is computed as the average of the pairwise Euclidean distances between each descriptor  $x_i^m$  and its K nearest neighbors  $\{x_{ij}^m\}_{j=1}^K$  (Sanchez-Martinez et al., 2017). The input descriptors we consider here are the HR (i.e.  $x_i^1 \in \mathbb{R}$ ) and the longitudinal velocity values along each cycle. As the dimension of the latter varies over cycles, all cycles were resampled along the temporal axis so that  $x_i^2 \in \mathbb{R}^{d_2}$  (we set  $d_2 = 65$ ).

Based on  $\{K_m\}_{m=1}^M$ , a set of sample-wise matrices  $\{\mathbb{K}^i\}_{i=1}^N$  is defined. Each  $\mathbb{K}^i$  encodes the similarity of sample *i* to the other samples taking into account the different descriptors. In practice,  $\mathbb{K}^i$  is built from stacking the  $i^{th}$  columns of all kernel matrices  $\{K_m\}_{m=1}^M$ . Formally, the projection of sample *i* is expressed as

$$y_i = A^T \mathbb{K}^i \beta \quad . \tag{3}$$

Plugging (3) into (1), the optimization problem becomes

$$\min_{A,\beta} \sum_{i,j} \|A^T \mathbb{K}^i \beta - A^T \mathbb{K}^j \beta\|^2 \mathbb{W}_{ij}$$
(4)

s.t. 
$$\sum_{i} \|A^T \mathbb{K}^i \beta\|^2 \mathbb{D}_{ii} = 1,$$
 (5)

$$\beta_m \ge 0, \sum_m \beta_m = 1 \tag{6}$$

where  $\mathbb{W}$  is the multiview generalization of W in (1), a global affinity matrix computed by combining all the individual kernel matrices (in this paper we used  $\mathbb{W} = \frac{1}{M} \sum_{m} K_{m}$ , with kernel matrices  $\{K_{m}\}_{m=1}^{M}$  being normalized across features prior to the summation through a variance-based method, described by Sanchez-Martinez et al. (2017)). The constraint in (5), with  $\mathbb{D}_{ii} = \sum_{j} \mathbb{W}_{ij}$ , removes an arbitrary scaling factor in the output embedding.

Minimizers  $A^*$  and  $\beta^*$  are obtained by an iterative two-step optimization strategy (Lin et al., 2011). At each iteration, A and  $\beta$  are alternately fixed to the value of last-step's solution and the problem is solved for the other. Iterations stop once a convergence criterion is met (e.g. maximum number of iterations or stable value of cost function). Solving (4) for A amounts to a generalized eigenvalue problem: the columns of the optimal A are the corresponding eigenvectors. Solving (4) for  $\beta$ , on the other hand, corresponds to a nonconvex quadratically constrained quadratic programming problem. To obtain a low-dimensional representation, one can choose the columns of A associated to the k lowest eigenvalues, yielding  $A \in \mathbb{R}^{N \times k}$  and thus  $y_i \in \mathbb{R}^k$ , i = 1, ..., N.

Once A and  $\beta$  have been learnt, the projections of the training samples can be computed using (3). Moreover, a new sample z can be mapped into the low-dimensional space by

$$y_z = A^T \mathbb{K}^z \beta \quad , \tag{7}$$

 $\mathbb{K}^z \in \mathbb{R}^{N \times M}$  and  $\mathbb{K}^z(n,m) = k_m(x_n,z)$ .

Thus, the projection of new samples is determined by the similarities of their input-space features with those of the samples in the training set.

## 2.4. Discriminative analysis and physiological interpretation

In the low-dimensional space, the spatial distribution of the projected cycles is learned in an unsupervised way, solely based on their input-space similarities and not taking into account any label (i.e. control/ANT1) information. Our aim is to explore this simplified representation towards the identification of distinctive clusters of response by the two populations, and the unraveling of the pathophysiological mechanisms behind such differences.

We perform two levels of analysis (see stage *III* of Figure 1): one that is based on the overall spatial distribution of samples of each population in the output space (i.e., not distinguishing subjects), and another where we cluster the subjects based on the trajectories defined by their sequences in the output space.

#### 2.4.1. Cycle-wise analysis: population signatures

To obtain the predominant patterns of response of each population, we i) draw a path passing through the regions of higher density of both healthy and diseased populations, and ii) sample the path at multiple points and adopt a multiscale adaptation of kernel regression (MKR) (Bermanis et al., 2013; Duchateau et al., 2013) to backproject them to input-space patterns. We hypothesize that analyzing the evolution of input features along this path will highlight discriminative characteristics of the diseased population.

Each such point q is backprojected based on an interpolation/regression from the known Y and X. A Gaussian kernel k is used to evaluate its similarity  $k(q, y_i)$  with each  $\{y_i\}_{i=1}^N \in Y$ ; its reconstruction in the space of feature m, here denoted as  $f_m(q)$ , is based on the known input-space representations  $X_m = \{x_i^m\}_{i=1}^N$  and weighted by such similarities:

$$f_m(q) = \sum_{i}^{N} k(q, y_i) b_{mi} \tag{8}$$

where  $b_{mi}$  stands for the  $i^{th}$  column of matrix

$$B_m = \left(K + \frac{1}{\gamma_m}I\right)^{-1}X_m\tag{9}$$

with  $K = [k(y_i, y_j)]$ ,  $\gamma_m$  a regularization weight and I the identity matrix. A multiscale approach is adopted where  $f_m$  is updated in an iterative coarse-to-fine process, with the kernel bandwidth halved at each step, from the maximum to the average output-space neighborhood size (details in (Duchateau et al., 2013)).

#### 2.4.2. Sequence-wise analysis: subjects' trajectories

The idea behind the trajectory-based analysis is that the trajectories defined by the projected cycles of each subject (in temporal order) can be considered physiological descriptors of response to stress, and, as such, performing cluster analysis in the trajectory space can help us identify how all the subjects are organized in groups of response.

For each subject p, the trajectory defined by the projected data consists of a multidimensional  $C_p \times k$  matrix, where  $C_p$  is the number of cycles of subject p's sequence. An element (c, dim), c = $1, ..., C_p$ , dim = 1, ..., k, tells us how the mode of variation associated with dimension dim is being expressed at cycle c. Intuitively, the whole trajectory matrix encodes a weighted combination of the k modes of variation at each cycle of the sequence. Our hypothesis is that there will be differences in the trajectory matrices of the two populations, specific to the ANT1 pathology.

To cluster trajectories, as each subject's sequence has a different length, and different ratios of baseline/stress/recovery durations, standard distance metrics cannot be applied. For that reason, the DTW algorithm is used. This algorithm allows aligning two multidimensional time series by stretching sections in the temporal axis (one-tomany correspondence) in such way that some distance metric (Euclidean in our case) between the aligned time series is minimized (Bemdt and Clifford, 1994). Prior to the DTW alignments, trajectories are slightly smoothed using total variation denoising (Rudin et al., 1992) (denoising weight  $\lambda$ = 0.01), to reduce noisy oscillations while preserving sharp transitions corresponding to state changes (rest-stress-recovery). Finally, a distance matrix is built from the pairwise distances and fed to a hierarchical clustering algorithm (Ward Jr., 1963), and the results are compared with the known labels.

#### 3. Experiments and Results

#### 3.1. Parameterization

Experiments were ran with several parameterizations. Alternatively to the standard iterative optimization process described in Section 2.3, having  $\beta \in \mathbb{R}^2$  and  $\sum_i \beta_i = 1$ ,  $\beta_i > 0$ , we simply performed a grid search on a discrete set of vectors obeying  $\beta = [\beta_1, 1 - \beta_1]^T$ ,  $0 < \beta_1 < 1$ , used them for initialization, and solved the corresponding generalized eigenvalue problem for the projection matrix A. In other words, we ran one single iteration of the standard optimization process for different initializations of the weight vector  $\beta$ .

Table 1 lists the parameterization corresponding to the results shown and discussed in this section. We denote by  $k_{\sigma}$  and  $k_{sparse}$  the number of neighbors used in the estimation of the kernel bandwidths and in a sparsing step of the global affinity matrix, respectively (refer to Sanchez-Martinez et al. (2017), Nogueira et al. (2017) for further details).

From our experiments, we found that the results presented a relatively low sensitivity to the values of  $\beta$ ,  $k_{sparse}$  (except for very small values), and higher sensitivity to the value of  $k_{\sigma}$ . Lower kernel bandwidths mean higher sensitivity to variations in the data, and vice-versa. We heuristically tuned the value of  $k_{\sigma}$  having in sight a good trade-off between the spread and the spatial smoothness of the output-space data distribution.

For the MKR, we decided to use the first 6 dimensions of the projected data, since including further dimensions had little influence in the reconstructed modes (higher dimensions encode more noisy variability). The value of  $\gamma_m$  in (9) was tuned to minimize the average curve reconstruction error over 150 fixed samples (10 of each subject).

Table 1: Parameterization details. Feature weight vector defined as  $\beta = [\beta_{HR}, \beta_{velocity}]^T$ .

Data	
N	1377
M	2
MKL	
$k_{\sigma}$	$0.05 \times N$
$k_{sparse}$	$0.25 \times N$
β	$[0.5, 0.5]^T$
MKR	
dimensionality $k$	6
$\gamma_m$	0.1

## 3.2. Population-wise analysis: representative signatures

We computed the low-dimensional representation of the data using MKL, with the parameterization in Table 1.

Figure 3 displays the probability density functions learned from the distributions of control (blue) and ANT1 patient (red) samples, considering the pairwise combinations of dimensions 1-4, obtained using the non-parametric method of kernel density estimation (Epanechnikov, 1969). We focused the analysis on the two first dimensions



Figure 3: Probability density function for the control and ANT1 sample distributions, considering pairs of the first dimensions of the projected data.



Figure 4: First two dimensions of projected data colored according to control (blue) and ANT1 (red) labeling (left). Separated distributions of control (middle) and ANT1 patient (right) samples, colormapped according to a stress score, consisting of the normalized HR sequence of each subject, mapped to the [0,1] interval.

(column 1), as including further dimensions did not add any additional insight from a physiological perspective (see in Appendix A an analogous analysis including dimension 3).

A scatter plot of the projected data, using the first two dimensions, is shown in Figure 4-left. In this plot, each point refers to a cardiac cycle of a subject, and is colored according to the control/ANT1 (blue/red) label. In the middle and right columns of the same Figure, we isolate each population's distribution of output-space samples and color them according to a stress score (computed as the HR value normalized by the minimal and maximal HR values of the corresponding patient). In both, the trend is to gradually transition from white (baseline/recovery) to dark blue/red (peak stress) in the counterclockwise direction. In fact, there is a continuum of response defined by the two distributions, where the ANT1 distribution is positively shifted in that same direction with respect to the control's.

To interpret the physiological implications of this



Figure 5: Distribution-based modes of variation. Left: Path over the distributions of control and ANT1 patients that were used for the estimation of the distribution-based modes of variation. The plotted points were the inputs for MKR. Right: MKR results, with color correspondence with the plotted path points in the left plot; the curves corresponding to the 4 annotated points are plotted with thicker linewidths.



Figure 6: Systole dynamics. Left: velocity curves of Figure 5-right, non-normalized for HR. Middle: corresponding displacement curves. A marker is plotted on the systolic peak for each curve. Right: analysis of the timing and total displacement at the systolic peak, throughout the path drawn in Figure 5-left.

spatial shift between distributions, we drew a path over the higher density regions and used MKR to reconstruct the velocity curves along such path. The path and points to backproject are shown in Figure 5-left, whereas the reconstructed mode of variation is shown in Figure 5-right.

Focusing on the evolution of the velocity curves from baseline to peak stress for the control population (see annotations in Figure 5-left), we observe a relative systolic lengthening, with the systolic peak happening later in the cycle, and a gradually shorter diastole, reaching some degree of E-A merging at peak stress. These are typical signatures of a normal response to exercise. Looking at the velocity patterns corresponding to the ANT1 baseline region (annotated in Figure 5), it is evident that ANT1 patients start off with some of these exercise signatures (e.g. shorter diastole); on the other hand, they also show a more accentuated augmentation of these signatures at peak stress (e.g. reaching complete E-A fusion), together with some additional shape changes in the velocity profile (especially noticeable during systole). Focusing on systolic function in particular, we integrated the velocity curves to obtain the corresponding displacement profiles (Figure 6-middle), and plotted the timing of end-systole against the corresponding displacement (Figure 6-right). In these plots, the timing is not normalized for HR. It is observed that, as HR goes up, controls increase the peak contraction and, while the absolute ejection time reduces, the relative duration of systole with respect to the cycle length increases; on the contrary, ANT1 patients fail to modulate timings of events and contractility in the same manner, as they reach peak stress. The findings are in agreement with the results of Sanchez-Martinez et al. (2017), that linked E-A fusion and reduction of contractility to exercise response in HFPEF, in a standard exercise context.

#### 3.3. Sequence-wise analysis: subjects' trajectories

The trajectory defined by the projected samples of each particular subject in the output space (Figure 5-left) carries information regarding how that subject responds in terms of the patterns recovered in Figure 5-right. Thus, we took such trajectories as descriptors of response to stress, and used them to cluster the subjects into groups of response. For that, we computed the pairwise distances among all subject trajectories (using DTW, as described in 2.4.2), and used them as inputs to a hierarchical clustering algorithm. Other than this hard clustering, we used multidimensional scaling (MDS) to visualize a 2D scatter plot representation of the subjects, based on the same pairwise distances. Good clustering results were obtained even considering only the first dimension of the trajectories – in Figure 7a, the two clusters are even linearly separable, although there is a very subtle transition between them (MDS distributions accounting for more dimensions can be found in Appendix A). In this scenario, the hierarchical clustering had perfect accuracy in the separation of healthy and diseased subjects (Figure 7b).

## 4. Discussion

We proposed an analysis framework for complex datasets composed of continuous multiview data sequences extracted from stress echo acquisitions, with the main objectives of i) discriminating healthy and pathological clusters of response and ii) understanding the underlying pathophysiological mechanisms. The framework extends the previous work by Sanchez-Martinez et al. (2017) to nonstandardized stress echocardiography. Transitioning from standardized to nonstandardized data implies that each subject is no longer represented by a single point in the MKL output space, but by a variable number of points (with an associated time order). The main contribution of the proposed framework lies in the concept of studying low-dimensional trajectories for clinical interpretation, an under-explored way to look at multiview clinical time series.

The discriminative power of the framework was first confirmed in Figure 3, where distinctive regions of higher control/ANT1 sample density were identified. It was again confirmed in Figures 7a and 7b, where a cluster analysis based on the subject trajectories in the output space accurately grouped them according to the diagnostic label.

Although there were distinctive regions of higher data density from the two populations, the two distributions were organized in a continuum of data. Such continuum was observed to be correlated with the stress level (Figure 4), and the ANT1 population data seemed to be positively shifted in the stress direction, when compared to the control distribution. In other words, the physiological patterns of the ANT1 population, at baseline conditions, are comparable to those found in controls during mild exercise.

To provide an idea of the variability found in trajectories (thus, signature intensities) within both control and ANT1 populations, we display the subject-wise projections in Figure C.12 in Appendix C.

While the studied populations used to illustrate the framework are distinctively different and easy to clinically discriminate based on heart failure symptoms at even the least exercise, our analysis can potentially provide novel insight in the physiology of this genetic mutation. However, the modest number of patients in the study precludes from any final conclusion when it comes to more in-depth pathophysiology analysis. On the other hand, the rarity of this mutation impedes the gathering of large datasets.

Among the main challenges of dealing with nonstandardized echocardiography sequences were those related to data processing and feature extraction, due to the complex nature of the data. Some recurrent problems that were especially likely to occur during stress were: noisy ECG, as in some cases the R peak became indiscernible and automated segmentation was not possible; out-of-plane heart motion resulting in an absence of Doppler signal, and significant breathing motion relative to the defined region of interest. However, the final mapping to the low-dimensional space was found to be fairly robust to outliers (i.e. the obtained modes of variation/projections were not overfitting the outliers, e.g. cycles with saturation peaks, cycles with



Figure 7: Trajectory clustering based on DTW distance matrix (for these results, only the first dimension of the trajectories was considered). (a) 2D scatter plot of the subjects with MDS. (b) Hierarchical clustering (subjects 0-9 correspond to controls and subjects 10-14 correspond to ANT1 patients).

no velocity signal, badly segmented cycles due to bad quality ECG regions – which also gave origin to unphysiological values of HR, etc.), so there was no need to perform a preselection of cardiac cycles based on signal quality. This would probably not be the case if we did not have a fairly large ( $\approx 1400$ samples) dataset. The poor-quality cycles could be in most cases recognized based on the projection values in the output space (e.g. an overall poorquality acquisition of subject 13 explains its considerable distance from the other subjects in the MDS plot (Figure 7a).

A high correlation between HR and the first dimension of the output space was found (0.81). To discard the possibility of HR strongly biasing the results, we repeated the whole experiments without feeding any HR information to the MKL algorithm (i.e. using only the velocity feature). After this, the correlation between HR and the first dimension of the output space remained high (0.67), with the main configuration of the two distributions and the corresponding modes of variation remaining similar (Appendix B). While the need of a multiview instead of a single-view dimensionality reduction algorithm could be arguable for this particular case, we still believe that taking HR as a feature to estimate the low-dimensional representation of the data can provide additional insight regarding the physiological interpretation and, in a scenario like ours, merging the two correlated features can add robustness when compared to pursuing a single-view approach on the velocity. Using more input features (e.g. velocity traces at other locations than the basal septum) would potentially allow a more specific characterization of the ANT1 response. Moreover, besides extending this analysis to velocity traces at other locations in the left and right heart chambers, one could also consider analyzing flow changes.

Despite the listed limitations, we demonstrated that the proposed framework was able to reveal the progression of the pattern of response from the control to the pathological domain. The size of the dataset would not have permitted to come up with this pattern by visual inspection of the data.

The proposed framework can be flexibly adapted to the study of any given pathology, keeping in mind that the definition of the relevant set of features should be, naturally, carefully thought in a pathology-dependent manner.

## 5. Conclusion

We have proposed a framework for the analysis of nonstandardized stress echocardiography sequence data. It uses unsupervised multiple kernel learning to merge myocardial velocity and heart rate information and obtain a low-dimensional representation of the data. The analysis is then performed in the new space, with multiscale kernel regression bridging the two spaces for interpretability. The framework is illustrated on handgrip exercise sequences acquired on a population of healthy controls and ANT1 mutation patients. The results show that the framework is able to detect distinctive clusters of response and provide insight into the underlying pathophysiological mechanisms, demonstrating its ability to handle this complex type of datasets, and the potential of nonstandardized protocols such as handgrip exercise for unmasking differential response mechanisms.

#### Acknowledgement

This work is supported by the European Union's Horizon 2020 Programme for Research and Innovation, under grant agreement No. 642676 (CardioFunXion), by the Fundació La Marató de TV3 (No. 20154031), and by the Spanish Ministry of Economy and Competitiveness under the María de Maeztu Units of Excellence Program (MDM-2015-0502). The authors thank Kevin Strauss MD and the Clinic for Special Children (Pennsylvania, USA), for providing access to the ANT1 mutation study data.

## Appendix A. Results for other combinations of output-space dimensions

Here, we show the sample-wise analysis when using dimensions 1 and 3 of the low-dimensional representation (Figure A.8a), and 2 and 3 (Figure A.9a). We also show the clustering of trajectories when using more than one dimension (Figure A.10, middle and right).

## Appendix B. Experiments without HR

In this appendix we show the results of the sample-wise analysis using only velocity data (i.e. no HR data) as input to the MKL framework (Figure B.11).

## Appendix C. Individual sequence lengths and projections

Herein, we detail individual sequence lengths and display the individual projections of the 15 subjects onto the 2d MKL space (Figure C.12).

#### References

- Bach, F.R., Lanckriet, G.R.G., Jordan, M.I., 2004. Multiple kernel learning, conic duality, and the smo algorithm, in: Proceedings of the Twenty-first International Conference on Machine Learning, ACM. pp. 6–. doi:10.1145/ 1015330.1015424.
- Bemdt, D.J., Clifford, J., 1994. Using dynamic time warping to find patterns in time series.
- Bermanis, A., Averbuch, A., R.Coifman, R., 2013. Multiscale data sampling and function extension. Applied and Computational Harmonic Analysis 34, 15–29. doi:10. 1016/j.acha.2012.03.002.
- Davidavicius, G., Kowalski, M., Williams, R., D'hooge, J., Salvo, G.D., Pierre-Justin, G., Claus, P., Rademakers, F., Herregods, M., Fraser, A., Pierard, L., Bijnens, B., Sutherland, G., 2003. Can regional strain and strain rate measurement be performed during both dobutamine and exercise echocardiography, and do regional deformation responses differ with different forms of stress testing? Journal of the American Society of Echocardiography 16, 299–308. doi:10.1016/S0894-7317(02)74428-3.
- Duchateau, N., Craene, M.D., Sitges, M., Caselles, V., 2013. Adaptation of multiscale function extension to inexact matching: Application to the mapping of individuals to a learnt manifold. Lecture Notes in Computer Science , 578–586.doi:\url{10.1007/978-3-642-40020-9\_64}.
- Epanechnikov, V.A., 1969. Non-parametric estimation of a multivariate probability density. Theory of Probability & Its Applications 14, 153–158. doi:10.1137/1114019.
- Hearst, M.A., Dumais, S.T., Osuna, E., Platt, J., Scholkopf, B., 1998. Support vector machines. IEEE Intelligent Systems and their Applications 13, 18–28. doi:10.1109/5254. 708428.
- Helfant, R.H., Devilla, M.A., Meister, S.G., 1971. Effect of sustained isometric handgrip exercise on left ventricular performance. Circulation 44, 982–993. doi:10.1161/01. CIR.44.6.982.
- Hotelling, H., 1936. Relations between two sets of variables. Biometrika 28, 321–377. doi:10.1093/biomet/28. 3-4.321.
- Kivowitz, C., Parmley, W.W., Donoso, R., Marcus, H., Ganz, W., Swan, H.J.C., 1971. Effects of isometric exercise on cardiac performance. Circulation 44, 994–1002. doi:10.1161/01.CIR.44.6.994.
- Li, Y., Yang, M., Zhang, Z.M., 2018. A survey of multi-view representation learning. IEEE Transactions on Knowledge and Data Engineering , 1–1doi:10.1109/TKDE.2018. 2872063.
- Lin, Y., Liu, T., CS, F., 2011. Multiple kernel learning for dimensionality reduction. IEEE Transactions on Pattern Analysis and Machine Intelligence 33, 1147–60. doi:10. 1109/TPAMI.2010.183.
- Nogueira, M., Piella, G., Sánchez Martínez, S., Langet, H., Saloux, E., Bijnens, B., De Craene, M., 2017. Characterizing patterns of response during mild stress-testing in continuous echocardiography recordings using a multiview dimensionality reduction technique, 502–513.
- Peressutti, D., Sinclair, M., Bai, W., Jackson, T., Ruijsink, J., Nordsletten, D., Asner, L., Hadjicharalambous, M., Rinaldi, C.A., Rueckert, D., King, A.P., 2017. A framework for combining a motion atlas with nonmotion information to learn clinically useful biomarkers: Application to cardiac resynchronisation therapy response prediction. Medical Image Analysis 35, 669



Figure A.8: Analysis plots for the combination of output-space dimensions 1 and 3. (a) Regions of baseline, stress and recovery for the two populations; (b) velocity patterns reconstructed from the distribution.



Figure A.9: Analysis plots for the combination of output-space dimensions 2 and 3. (a) Regions of baseline, stress and recovery for the two populations; (b) velocity patterns reconstructed from the distribution.

- 684. URL: http://www.sciencedirect.com/science/ article/pii/S1361841516301815, doi:https://doi.org/ 10.1016/j.media.2016.10.002.
Puyol-Antón, E., Ruijsink, B., Gerber, B., Amzulescu, M.S., Langet, H., De Craene, M., Schnabel, J.A., Piro, P., King, A.P., 2019. Regional multi-view learning for cardiac motion analysis: Application to identification of dilated cardiomyopathy patients. IEEE Transactions on Biomed-



Figure A.10: Changes in the MDS plots when considering more dimensions of the subject trajectories.



Figure B.11: Experiments without considering HR as input feature of MKL: projected data and reconstructed mode of variation of the velocity feature.



Figure C.12: 2D MKL projection of the full dataset (dark green), superimposed by each subject's individual projection, colormapped based on label (blue - controls; red - ANT1) and color-coded based on stress score, as previously done in Figure 4. Subject sequence lengths (from top left to bottom right): 74, 143, 64, 118, 64, 89, 63, 102, 67, 58, 95, 88, 110, 137, 105.

ical Engineering 66, 956–966. doi:10.1109/TBME.2018. 2865669.

- Rudin, L.I., Osher, S., Fatemi, E., 1992. Nonlinear total variation based noise removal algorithms. Physica D: Nonlinear Phenomena 60, 259 – 268. doi:10.1016/ 0167-2789(92)90242-F.
- Sanchez-Martinez, S., Duchateau, N., Erdei, T., Fraser, A.G., H.Bijnens, B., Piella, G., 2017. Characterization of myocardial motion patterns by unsupervised multiple kernel learning. Medical Image Analysis 35, 70–82. doi:10.1016/j.media.2016.06.007.
- Sanchez-Martinez, S., Duchateau, N., Erdei, T., Kunszt, G., Aakhus, S., Degiovanni, A., Marino, P., Carluccio, E., Piella, G., Fraser, A.G., Bijnens, B.H., 2018. Machine learning analysis of left ventricular function to characterize heart failure with preserved ejection fraction. Circulation: Cardiovascular Imaging 11. doi:10.1161/ CIRCIMAGING.117.007138.
- Strauss, K.A., DuBiner, L., Simon, M., Zaragoza, M., Sengupta, P.P., Li, P., Narula, N., Dreike, S., Platt, J., Procaccio, V., Ortiz-González, X.R., Puffenberger, E.G., Kelley, R.I., Morton, D.H., Narula, J., Wallace, D.C., 2013. Severity of cardiomyopathy associated with adenine nucleotide translocator-1 deficiency correlates with mtdna haplogroup. Proceedings of the National Academy of Sciences 110, 3453–3458. doi:10.1073/pnas.1300690110.
- Velasco, M., Gómez, J., Blanco, M., Rodriguez, I., 1997. The cold pressor test: pharmacological and therapeutic aspects. American Journal of Therapeutics 4, 34–38. doi:10.1097/00045391-199701000-00008.
- Voigt, J.U., Exner, B., Schmiedehausen, K., Huchzermeyer, C., Reulbach, U., Nixdorff, U., Platsch, G., Kuwert, T., Daniel, W.G., Flachskampf, F.A., 2003. Strain-rate imaging during dobutamine stress echocardiography provides objective evidence of inducible ischemia. Circulation 107, 2120-2126. doi:10.1161/01.cir.0000065249.69988.aa.
- Ward Jr., J.H., 1963. Hierarchical grouping to optimize an objective function. Journal of the American Statistical Association 58, 236–244. doi:10.1080/01621459.1963. 10500845.
- Wold, H., 1985. Partial Least Squares. John Wiley & Sons, Inc. doi:10.1002/0471667196.ess1914.pub2.
- Xu, C., Tao, D., Xu, C., 2013. A survey on multi-view learning. arXiv:1304.5634.