



Published in final edited form as:

IEEE Trans Biomed Eng. 2018 April ; 65(4): 733–744. doi:10.1109/TBME.2017.2716365.

eCurves: A Temporal Shape Encoding

Elena Bernardis,

Department of Pediatrics, Children's Hospital of Philadelphia, Philadelphia, PA 19104, USA

Yong Zhang,

Colin Artificial Intelligence Lab, Richmond, British Columbia, Canada

Ender Konukoglu,

ETH Zurich, Zurich, Switzerland

Yangming Ou,

Department of Medicine and Radiology, Boston Children's Hospital, Harvard Medical School, Boston, MA 02215, USA

Harold S. Javitz,

Center for Health Sciences, SRI International, Menlo Park, CA 94025, USA

Leon Axel,

Department of Radiology, NYU School of Medicine, New York, NY 10016, USA

Dimitris Metaxas,

Department of Computer Science, Rutgers University, Piscataway, NJ, 08854 USA

Benoit Desjardins, and

Department of Radiology, University of Pennsylvania, PA 19104, USA

Kilian M. Pohl

Center for Health Sciences, SRI International, Menlo Park, CA 94025, USA

Abstract

Objective—This paper presents a framework for temporal shape analysis to capture the shape and changes of anatomical structures from 3D+t(ime) medical scans.

Method—We first encode the shape of a structure at each time point with the spectral signature, *i.e.*, the eigenvalues and eigenfunctions of the Laplace operator. We then expand it to capture morphing shapes by tracking the eigenmodes across time according to the similarity of their eigenfunctions. The similarity metric is motivated by the fact that small shape deformations lead to minor changes in the eigenfunctions. Following each eigenmode from the beginning to end results in a set of eigenmode curves (eCurves) representing the shape and its changes over time.

Results—We apply our encoding to a cardiac data set consisting of series of segmentations outlining the right and left ventricles over time. We measure the accuracy of our encoding by

However, permission to use this material for any other purposes must be obtained from the IEEE by sending an email to pubs-permissions@ieee.org.

Corresponding author: Kilian M. Pohl (kilian.pohl@sri.com).

training classifiers on discriminating healthy adults from patients that received reconstructive surgery for Tetralogy of Fallot (TOF). The classifiers based on our encoding significantly surpasses deformation based encodings of the right ventricle, the structure most impacted by TOF.

Conclusion—The strength of our framework lies in its simplicity: it only assumes pose invariance within a time series but does not assume point-to-point correspondence across time series or a (statistical or physical) model. In addition, it is easy to implement and only depends on a single parameter, *i.e.*, the number of curves.

I. Introduction

The diagnoses and further understanding of medical diseases often rely on accurately describing structural changes of anatomy as captured by 3D+t(ime) medical scans. For example, Figure 1 visualizes the deformation of ventricles across one heart beat as a series of 3D models generated from segmentations of cine MRIs. In the medical imaging domain, these changes are frequently quantified based on low-dimensional, volumetric measurements, such as ejection fraction [1], [2] or regression based on the volumes of anatomical structures [3], [4], or high-dimensional, deformation-based encodings such as [5], [6], [7]. We start by reviewing existing temporal shape encodings and then propose an alternative by expanding a low-dimensional 3D shape encoding to the time domain.

Compared to volumetric encodings, representations based on deformation maps are considered more sensitive towards capturing shape changes of anatomy [8], [9], [10]. These encodings non-rigidly register the individual images (*i.e.*, time points) and corresponding segmentations of a 3D+t scan so that the resulting deformation maps capture local expansion or contraction of anatomical structures between time points. They have been heavily applied to anatomical shapes extracted from cardiac MRIs [11], [12], [13]. However, the registration of these encodings can bias measurements as they are complex systems that are based on simplifying assumptions and expert-based parameter tuning [14]. For disease modeling based on the corresponding deformation maps, one furthermore needs to account for the ‘Curse of Dimensionality’ as the maps are generally high dimensional and describe changes only between adjacent time points for each voxel. The computer vision community addresses the Curse of Dimensionality by embedding dimensionality reduction methods within temporal shape analysis. In the context of human action recognition, shape changes are encoded by first generating large sets of features by, for example, matching objects across space and time [15], [16], [17], [18] or solving Poisson equations [19]. These maps are then reduced to a compact set of ‘shape-time features’ implicitly capturing local properties of the changing shapes [19], [20], [21], [22], [23] by approaches such as dictionary learning [20].

In this paper, we propose an alternative low-dimensional representation, which encodes shape changes with a relatively small number of features to start with. Thus, our encoding minimizes the need for dimensionality reduction methods, which should only be used on data that complies with their underlying assumptions. Based on the spectral signature [24], [25], our proposed encoding describes a 3D shape by the eigensystems of the Laplace operator [26], [27], [28] applied to the corresponding binary map. The eigenvalues of that

system implicitly carry information about the shape, such as curvature, surface area and volume. Their popularity in the vision and graphics community [29], [30], [31] is partly due to their versatility as prior statistical or physical modeling is not required. Furthermore, the representation is pose invariant [32] so that using eigensystems to compare several shapes does not require prior registration of the objects to a common coordinate system. Finally, small shape deformations lead to minor refinements in the eigenvalues (see [28] for a theoretical discussion and [33] for empirical examples). We use this property to extend spectral signatures to the temporal domain and thus create a low dimensional encoding of changing shapes.

An overview of our temporal shape representation is given in Figure 2. First, we compute the spectral signature for each time point, yielding a family of temporal eigenvalues or **eValues**. We then track these eigenmodes (**eTracking**) across time points, to ensure that the eigenvalues measured at different time points are properly linked. Following a mode from the first to the last time point results in an *eigenmode Curve* (**eCurve**). The set of eCurves across all modes then encodes the temporal shape changes.

The strength of our representation lies in its simplicity. By defining our temporal shape representation based on the spectral signature, we assume no spatial correspondence across time series (e.g., the beating hearts of different subjects), or prior physical or statistical models of the data. Thus, comparing curves across subjects simply reduces to comparing eigenvalues across time points (see [33] for an example). Within a time series, however, the corresponding eigenfunctions are linearly aligned to a common coordinate system so that they (and consequently the eigenvalues) can be compared across time points. In this sense, tracking is similar to approaches that match shapes based on finding correspondences between spectral signatures, such as [34], [35], [36]. For the sake of simplicity, our tracking approach minimizes algorithmic and computational complexity by assuming minor shape changes between time points. This rather simple model is supported by experiments. Finally, our representation only depends on a single parameter, the number of curves, for which we provide a decision criteria. Given these constraints, eCurves is relevant to a wide range of applications in medical image analysis.

This paper is an extension of our work in [37]. The main novelty is the introduction of eTracking: dropping the assumption that the order of the eigenmodes does not change across time, *i.e.*, constructing eCurves does not simply follow the eigenvalue of a mode across time but instead carefully matches them. As we show on a new synthetic example in Section II, the changes of eigenmodes across time do not necessarily relate to the changes in the deforming shape. We capture these changes by first tracking them across modes and time via eTracking and then constructing the curves from the tracked modes. In addition to introducing eTracking, we derive a criteria for setting the dimension of the eCurves in Section III. As in [37], we apply our encoding to a cardiac data set consisting of healthy controls and adult patients that received reconstructive surgery for Tetralogy of Fallot (TOF) during infancy. However, we generalize the results by increasing the data set from 20 to 88 hearts, replacing the leave-one-out validation with nested two-fold cross-validation, and analyze the accuracy of not only one but three different classifiers. To further improve the scientific rigor of the experiments, we omit methods reducing the dimensionality of the

curves and instead directly train the classifiers on the curves of volumetric, deformation-based, and our encoding.

We choose a data set consisting of TOF and healthy control subjects as a testing environment for our encoding for several reasons. First, the diagnosis of the subjects provides ground truth for the classifier as only TOF patients received reconstructive surgery during infancy. In addition, the task of identifying abnormal shape motion in this population is nontrivial. While reconstructive surgery during infancy often leads to abnormal ventricular shapes that are quite different from healthy hearts, *i.e.*, the patch created at the right ventricular outflow tract results in expansion of that area, the impact of reconstructive surgery in some cases is so subtle (dilation of the right ventricular combined with a flattening of the interventricular septum during diastole) that the ventricular shape motion of TOF patients is indiscernible from healthy hearts based on volumetric scores, the clinical standard for quantifying cine MRIs. The third and final reason for choosing this data set is that the clinical literature fails to define clear guidelines regarding the need of surgical intervention on the pulmonary valve based on the severity level of regurgitation within hearts of patients that had reconstructive surgery for TOF [38]. Improving measurements of the shape and function of patients' beating hearts, such as through our proposed representation, may allow clinicians to better quantify the long-term impact of this reconstructive surgery and thus better the decision process regarding follow-up surgery during adolescence or adulthood.

The paper is organized as follows: in Section II, we introduce the Laplacian operator and extend it to encode temporal changes. In Section III, we present the experiments, in which we show that eCurves is significantly better than deformation-based encodings of the right ventricle, the structure most relevant to TOF. In Section V, we end the article highlighting the potential of our representation and the importance of accurately capturing shape changes.

II. Temporal Shape Encoding

We now describe our representation for capturing a 3D shape and its changes across time. Specifically, Section II-A provides a brief introduction to the Laplace operator, which we use to compute the signatures of the 3D shapes at each time point (see also Figure 2). We then extend the signature to the temporal domain in Section II-B by ensuring consistency of the shape encoding across adjacent time points. In essence, we link the eigenmodes at adjacent time points by measuring similarity between the corresponding eigenfunctions. In Section II-C, we follow those links from the first to the last time point to create eCurves.

A. eValues: Capturing Shapes via the Spectrum of Laplace Operator

We represent each shape in the time series via the eigenvalues and eigenfunctions of the Laplace operator [37]. We start with a brief overview of the Laplace operator (see [28], [24] for a more thorough discussion) and then outline how to use this encoding to represent the series of shapes.

Let the segmentation of each anatomical structure be viewed as a closed bounded domain $\Omega \subset \mathbb{R}^d$ with piece-wise smooth boundaries. The Laplace operator Δ_Ω for a twice differentiable function f on Ω is then defined as

$$\Delta_\Omega f \triangleq \sum_{i=1}^d \frac{\partial^2}{\partial x_i^2} f, \quad (1)$$

with $x \triangleq \{x_1, \dots, x_d\}$ spatial coordinates. We can use this operator to describe shapes, whose boundaries are denoted by $\partial\Omega$, by solving the Helmholtz equation with Dirichlet type boundary conditions,

$$\Delta_\Omega f + \lambda f = 0, \quad \forall x \in \Omega \text{ and } f(x) = 0, \quad \forall x \in \partial\Omega, \quad (2)$$

where $\lambda \in \mathbb{R}$. A shape Ω is then encoded by the infinite series of eigenpairs $\{(\lambda_j, f_j)\}_{j=1}^\infty$ satisfying Equation 2. Note, there are many alternatives to computing the Laplace operator, such as omitting the need for boundary conditions by solving the Helmholtz equation with respect to a mesh [24], [39], [40], [35]. Regardless of the specific model, the diverging sequence of non-trivial eigenvalues $0 < \lambda_1 \leq \lambda_2 \leq \dots$ is called the spectrum of Ω .

Just using the eigenvalues alone results in a shape representation that has many desirable properties for medical imaging studies. First, they are invariant to isometric transformations and this fact can also be extended to scaling [24]. This property is very important for the analysis of shape series, as one can eliminate pose correction of the individual shapes, a frequent source of noise in 4D analysis. Second, the *heat-trace*

$$Z(\gamma) \triangleq \sum_{j=1}^{\infty} e^{-\lambda_j \gamma} = \sum_{m=0}^{\infty} a_{m/2} \gamma^{-d/2+m/2} \text{ with } \gamma > 0 \quad (3)$$

links the eigenvalues $\{\lambda_j\}_{j=1}^\infty$ to an object's geometry as the coefficients $a_{m/2}$ are related to geometric properties of Ω , such as its volume or its surface mean curvature [41]. Furthermore, eigenvalues and eigenfunctions associated with the lower modes are influenced by the overall shape of an object, while they are more influenced by regional shape fluctuations at higher modes. Finally, the eigenvalues change continuously when infinitesimal small deformations are applied to the object's boundary (see [33] for examples).

The property of continuity is an important factor for choosing the Laplacian spectrum as the basis of our temporal shape encoding. The aim of our representation is to accurately encode objects that continuously change their shape over time. Let $\Omega(t)$ be the shape of an object at time t and $\Lambda(t) \triangleq \{\lambda_1(t), \lambda_2(t), \dots\}$ be the corresponding spectrum, with respective eigenfunctions $F(t) \triangleq \{f_1(t), f_2(t), \dots\}$. The series of $\Lambda(\cdot)$ is therefore able to capture both

geometry and change of an object $\Omega(\cdot)$. In other words, we can use $\Lambda(\cdot)$ to encode the changes of a shape between time points.

Our goal is to encode a series of 3D binary segmentations outlining the same anatomy of a subject sampled at T different time points. We compute the first N eigenvalues and eigenfunctions for each of the segmentations. This results in a series of eigenpairs $(\Lambda(t), F(t)) \triangleq \{(\lambda_1(t), f_1(t)), (\lambda_2(t), f_2(t)), \dots\}$, *i.e.*, pairs of eigenvalues and eigenfunctions that encode information of the object's shape at specific time points. Figure 3a shows the eigenseries $\{(\Lambda(1), F(1)), \dots, (\Lambda(10), F(10))\}$ capturing the changes of an asymmetric shape (a torus with an additional small bump) at 10 different time points. The figure shows eigenfunctions $F(t)$ (top) and corresponding eigenvalues $\Lambda(t)$ (bottom) at modes $j = 1, 5, 8, 17, 20$ (y-axis) with a line connecting eigenvalues associated with the same mode to visualize the changes over time (x-axis). Inline with the mathematical proof of [33], the eigenfunctions nicely display how lower order modes capture coarser shape information through a few sinks and sources shown in dark red and blue, while higher modes contain many more sinks and sources making them sensitive to local fluctuations along the boundary, *i.e.*, the increase in the radius of the torus' tube over time.

B. eTracking: Tracking Eigenmodes across Time-Points

To correctly encode shape changes, we have to ensure that the eigenpairs between adjacent time points are properly correlated, as mode j at time point t does not necessarily have to correspond to mode j at time point $t + 1$. For example, the eigenfunctions in Figure 3a do not properly reflect the small deformations of the torus between adjacent time points, as eigenfunctions of modes 17 and 20 drastically change over time. We resolve this issue via eTracking: first compare eigenpairs at adjacent time points and then link the most similar ones.

eTracking is based on the property that the eigenfunctions of the Laplacian are orthonormal to each other, *i.e.*, given two eigenfunctions of $\Omega(t)$, $f_a(t)$ and $f_b(t)$, then

$$|\langle f_a(t), f_a(t) \rangle| = 1 \text{ and } \langle f_a(t), f_b(t) \rangle = 0 \text{ for } a \neq b, \quad (4)$$

where $|\cdot|$ is the absolute value and $\langle \cdot, \cdot \rangle$ is the dot product between the eigenfunctions. Note, the absolute value of the inner product accounts for the ambiguity in sign of eigenfunctions, *i.e.*, $f_a(t)$ and $-f_a(t)$ are both valid eigenfunctions.

Assuming the time series does not contain any pose changes and the shape does not change between two time points (*i.e.*, $\Omega(t) = \Omega(t + 1)$), then the same property would hold for all pairs of eigenfunctions across time points:

$$|\langle f_a(t), f_a(t+1) \rangle| = 1 \text{ and } \langle f_a(t), f_b(t+1) \rangle = 0 \text{ for } a \neq b, \quad (5)$$

where the product is defined by the union of the domains defined by the two eigenfunctions. Thus, one way to track eigenpairs between adjacent time points is according to their degree

of orthogonality, *i.e.* one minus the absolute value of the inner product between the corresponding eigenfunctions. A mode at time t is then linked to the one at time $t+1$ with whom it has the lowest degree of orthogonality or, equivalently, the highest absolute inner product. Now, let $\phi_a(t)$ represent the mode at time $t+1$ to which mode a at time t is linked to, *i.e.*, the mode at time $t+1$ with whom mode a at time t has the highest absolute inner product:

$$\phi_a(t) \triangleq \operatorname{argmax}_{b>0} |\langle f_a(t), f_b(t+1) \rangle|. \quad (6)$$

As we assume small shape deformations between time points and due to the orthonormal properties of eigenfunctions of the Laplacian, then the absolute inner product between linked modes will not only be much larger than the absolute inner product of Mode ‘ a ’ to the other modes at time $t+1$, *i.e.*,

$$|\langle f_a(t), f_{\phi_a(t)}(t+1) \rangle| >> |\langle f_a(t), f_c(t+1) \rangle| \forall c \neq \phi_a(t), \quad (7)$$

but the absolute inner product of the linked modes should also be close to one, *i.e.*,

$$|\langle f_a(t), f_{\phi_a(t)}(t+1) \rangle| \approx 1. \quad (8)$$

The robustness of this tracking model depends on the signal-to-noise ratio (SNR) with lower order modes being generally more robust than higher order modes (given the same SNR). To gain a better understanding of this property, Figure 4 shows on the left a torus with a “bump of interest” (reference), whose inner product (see plot at the bottom of Figure 4) is computed with respect to the tori on the right (*i.e.*, the shape at the next time point) with a larger bump of interest and two additional bumps. The size difference in the bump of interest between the reference and the shapes to the right represents the actual shape changes to be tracked, or the ‘signal level’ in our tracking model. The size of the two additional bumps relates to the “noise level”. As the additional bumps become larger, the SNR decreases and the tracking is less reliable as the noise causes a reduction in the inner product especially for higher order modes. As previously mentioned, higher order eigenfunctions carry information on finer shape details while lower order ones carry information on coarser shape details. Naturally, a reduction in SNR causes the tracking of finer shape details (and corresponding temporal changes) to become harder.

Very low SNR can also reduce the reliability of tracking lower order modes. For example, if the two additional bumps are equal or larger in size than the shape of the bump of interest (not shown) then the eigenfunction of the first mode would contain multiple sinks or sources. This would result in a low inner product with respect to the corresponding eigenfunction of the reference shape, which contains a single sink. In this low SNR scenario, the inner product would capture the influence of noise rather than the actual changes of the shape. Thus, the tracking of the modes would be unreliable across time. To account for the reliability issue with respect to the SNR, we associate proper linkage (or orthonormality) by

defining a minimum threshold τ for the absolute inner product $|\langle f_a(t), f_{\phi_a(t)}(t+1) \rangle|$. Modes that do not meet this threshold are omitted from our representation to minimize the risk of corrupting the proceeding analysis. Since the threshold is application dependent, we provide a criteria for setting τ on real data in Section III.

We complete the definition of eTracking by introducing $\Phi(t) \triangleq \{\phi_1(t), \phi_2(t), \dots\}$ to capture the links of all the modes between time t and $t+1$. If we let $(\lambda_a(t), f_a(t), \phi_a(t))$ denote the *eigen triplet* of mode a at time t , then eTracking at time t is defined as the eigen triplets across all modes:

$$(\Lambda(t), F(t), \Phi(t)) \triangleq \{(\lambda_1(t), f_1(t), \phi_1(t)), (\lambda_2(t), f_2(t), \phi_2(t)), \dots\}. \quad (9)$$

Note, that the order of the eigen triplets is defined by the ascending order of the eigenvalues at the first time point, *i.e.*, $\lambda_1(1) \leq \lambda_2(1) \leq \dots$. While in theory, the eigenvalues of two modes can be the same, *i.e.*, $\lambda_i(1) = \lambda_{i+1}(1)$, in practice we have not encounter this scenario so that the order across the triplets was unambiguous.

Returning to the deforming torus of Figure 3, the right side (b) shows the eTracked eigenvalues and eigenfunctions of the left side (a) according to ϕ . For the lower modes (1,4,8), the tracking did not change the order of the eigenfunctions and corresponding eigenvalues (see Figure 3a) as the eigenfunctions capture global changes in shape that are unaffected by the change in size of the torus and thus stable across time. The eigenfunctions of modes 17 and 20 are sensitive to the variations yielding drastic changes between adjacent time points when tracking is ignored. Organizing the modes based on eTracking, the corresponding eigenfunctions smoothly vary across time which results in the absolute values of the inner products of eigenfunctions of adjacent time points being close to 1. For example, the mean and standard deviation of the inner-products defined by the original order of Mode 20, *i.e.*, Figure 3a), is 0.82 ± 0.33 while for the eTracked series it is 0.94 ± 0.08 . The higher inner product is also reflected in the tracked eigenfunction, which now, like the modifications to the torus, gradually change over time. The corresponding series of the eigenvalues thus more accurately encode the changes of the torus across time, which, for the higher modes, are more distinct over time compared to the original series of eigenvalues.

C. eCurves: Capturing the Eigenseries as Curves across Time

Instead of just analyzing neighboring time points, we now construct eCurves, which capture both shape and its changes across all time points. For each mode a at the first time point, we define the eCurve as the set of eigenvalues that are linked according to $\Phi(t)$,

$$\Lambda^{(a)} \triangleq [\lambda_a(1), \lambda_{\phi'_a(1)}(2), \dots, \lambda_{\phi'_a(T-1)}(T)] \quad (10)$$

With $\phi'_a(1) \triangleq \phi_a(1)$ and for $1 < t < T$

$$\phi'_a(t) \triangleq \phi'_b(t) \text{ with } b \triangleq \phi'_a(t-1).$$

To comply with previous notation, $\Lambda^{(a)}(t)$ refers to the t^{th} entry of $\Lambda^{(a)}$. An important feature of this encoding is that it allows comparison of eCurves *across* time series without requiring any alignment between the time series. eCurves inherits this property from the eValues, which are pose invariant.

In summary, we encode temporal shape changes by first computing the eValues at each time point, tracking these values across adjacent time points (eTracking) and, finally, composing the eCurves from the resulting curves.

III. Experiments

We analyze the accuracy of our temporal shape encoding with respect to the right and left ventricles across one heart cycle of healthy and diseased subjects. The heart cycle of each subject is captured by segmentations of the ventricles at 30 time points (see also sample series in Figure 1). We next infer the accuracy of eCurves from that of three different automatic classifier. Specifically, we apply eCurves to the data set, feed the encodings into the classifiers and then record the sensitivity, specificity, and the accuracy of the classifier in discriminating healthy versus diseased subjects. Comparing these scores to the ones obtained from alternative temporal shape descriptors highlights the strength of eCurves. We describe our experimental setup in Section III-A and analyze the results in further detail in Sections III-B.

A. Experimental Setup

In our experiments, we evaluate eCurves and alternative temporal shape encodings on the cardiac data set. We review the experimental setup, in particular details on data acquisition and implementations of the temporal shape encodings, and the training and testing of the classifier via two-fold crossvalidation.

1) Data Set—The *cardiac data set* consists of segmentations of the left and right ventricle (such as in Figure 1) extracted from short-axis cine MRI scans of 42 healthy controls and 46 adult patients that received reconstructive surgery of Tetralogy of Fallot (TOF) during infancy. The cine MRI are acquired via a sequence of balanced steady state free precession short-axis images of the heart using breath holds in a 1.5 Tesla Siemens Avanto scanner with the following setting: 30 ms TR, 1.3 ms TE, 360×270 mm FOV; 8 mm slice thickness, 2 mm skip, 60–90 degree flip, 192×192 resolution and 1.25 mm in-plane resolution. We confine the spatial positions of these short-axis scans to the mid-portion of the ventricles (12 slices), where the pathologies are usually prominent. Additional details about the image acquisition of the cine MRI are described in [42]. [42] also outlines the procedure for segmenting the MRI at the enddiastole into the left ventricle (LV), consisting of myocardium and blood pool, and into the blood pool of the right ventricle (RV). Each label map of the end-diastole is propagated to the other time points via non-rigid registration [43]. Holes and

islands due to registration errors are removed by semi-automatically dilating and eroding the aligned registrations.

2) Computing Temporal Shape Encodings—We now describe in further detail the implementations of eCurves following the work flow outlined in Figure 2, as well as alternative encodings that we will use for comparisons: eValues, a deformation-based encoding, and a volumetric score.

In all of our experiments, a subject is represented by a series of segmentations over time that are rigidly aligned with each other. For each segmentation of a specific time point, we first compute a large number of eigenvalues and eigenfunctions (*e.g.*, 100), that solve the Helmholtz equation with Dirichlet type boundary conditions (Equation 2). Using the implementation of [33], we compute the Laplacian operator on the grid native to the image. Thus, we avoid the need of interpolation or mesh generation and work only with the information available in the images. Assuming a series consists of T time points, the anatomy of the subject across time is then represented by 100 T eigenvalues (**eValues**).

Since the eCurves proposed in Section II explicitly account for temporal constraints by tracking the modes across time, we compute the inner product between the eigenfunctions of the modes of adjacent time points and then link those with the highest inner product. While in theory each mode is uniquely linked with a mode at the next time point (unique match), in practice, the highest inner product of several modes may match with the same mode at the next time point (multiple matches). For example, Figure 5 shows two examples (left and right) of modes of the RV, whose highest inner product across all modes at the next time point matches with the same mode. While in the first example (left) Mode 12 at time point 15 will be linked to Mode 12 at time point 16, in the second example Mode 12 at time point 21 will be linked to Mode 13 at the next time point as its inner product is higher than that of Mode 13 at time point 21. To better understand the phenomena of multiple matches, we plot all modes across all 88 subjects (see Figure 6, first row) with the coordinate of each mode defined by its highest (x-coordinate) and second highest (y-coordinate) inner product across all 100 modes of the previous time point. *Uniquely matched* modes are plotted in orange while the blue dots depict *multiple matched* modes, *i.e.*, the presence of more modes at the previous time point whose highest inner product links to the same mode. The distribution of uniquely and multiple matched modes is very similar across the two structures, specifically 3.52% for LV and 6.58% for RV. As expected, the highest concentration of uniquely matched modes has the highest inner product close to 1 and second highest inner product close to zero. Interestingly enough, the difference between the highest and second highest inner product does not seem to predict if a mode is uniquely matched, as there are many uniquely matched modes with similar highest and second highest inner product. As multiple matches occur only in less than 7% across all modes and subjects, we reliably resolve the multiple match issue by applying the Gale-Shapley matching algorithm [44] to the inner products across all modes. The algorithm does so by recording, for each mode at time point ‘t-1’, the mode at ‘t’ yielding the highest inner product. Two scenarios can then occur: (a) The mode at time point ‘t’ temporarily accepts the possible link if the inner product is higher than that one of the mode it is currently linked to at time point ‘t-1’ (*e.g.*, Mode 12 in the example to the left of Figure 5), which is then annulled (called “trade up”). (b) If the link is

not accepted (*e.g.*, Mode 10 in the example on the left of Figure 5), then for the mode at time point ‘t-1’ the mode at time point ‘t’ is determined, which has the highest inner product of those the mode at time point ‘t-1’ has not been linked to yet. The process between (a) and (b) repeats until all modes at ‘t-1’ are temporarily linked to a mode at time point ‘t’. Note, other methods than the Gale-Shapley approach [44] exist for matching modes, such as the computationally more complex Hungarian algorithm used in [34], [40].

Having tracked the modes across time, we compute the eCurves for each case according to Equation 10 only keeping the first N curves, where N is set in correspondence to the ‘robustness threshold’ τ via a data-driven approach. To review, we only keep a subset of curves as we have previously shown [33] that higher modes correspond to progressively finer details of the shape that are not necessarily informative for distinguishing cohorts. Defining the subset by confining tracking to the first 25 modes, the issue of multiple matches is reduced according to Figure 6. Multiple matches only arises in 0.3% of matches associated with the LV and 0.9% of matches associated with the RV. This observation is in line with the discussion of Figure 4, which provides an example for the robustness of tracking of higher order modes being more sensitive to noise than the tracking of lower ordered ones. To insure the analysis not being corrupted by noise, eCurves omits all curves whose robustness scores are below the threshold τ . We then measure the robustness of curves by computing the 5th percentile score of the inner products between their tracked modes. Note, we choose the 5th percentile score over the overall minimum inner score so that the robustness measure is not influenced by noisy outliers. Setting τ via parameter exploration in combination with nested cross-validation (see Section III-A3 for details) then infers the setting of the optimal N , which is the highest order mode for which the curves associated with that and lower order modes receive robustness measures across all (training) subjects that are above τ . For example, Figure 7 plots the robustness measures for the first 100 modes across all samples. If τ would be set to 50%, the optimal N would then be 15.

We complete the construction of eCurves by sampling each curve associated with the ventricles at $T = 30$ time points ensuring that each subject is represented by the same temporal resolution. The scans were acquired in the clinical setting and thus vary with respect to the number of time points covering a heart beat (starting and ending at end-diastole). Curves constructed from time series with more or less than 30 time points are represented in the continuous domain through linear interpolation between time points. Their values at the 30 equally spaced time points are then recorded to ensure a uniform temporal resolution across all scans.

Figure 8 shows the eCurves for six subjects (3 healthy and 3 TOF), visualizing for each subject 3D models of the right (red) and left (green) ventricles at three different time points. The curves cover an entire cardiac cycle starting at the enddiastole and peaking at the end-systole. In comparison to the 3 healthy hearts, the peaks in the eCurves of the 3 TOF hearts seem less prominent in the right ventricle.

In addition to implementations based on eValues, we encode the shape changes of the ventricles via a deformation-based representation [42]. For each case and time point, we first reduce each image to the cardiac structures based on the segmentations of the RV and LV by

masking the rest of the image. We non-rigidly register the LV of the resulting image to the first time point of the time series via the Diffeomorphic Demons algorithm [45] and repeat the process for the RV. We compute the Jacobian Determinant from the resulting deformation map. We separate the LV of the first time point into 16 smaller regions based on the Bullseye plot suggested by the American Heart Association [46] and apply 66% of the bullseye plot to the RV resulting in 11 regions. For each subregion and time point, we compute the average of the Jacobian Determinant over that region. For each subject and subregion, we then interpret the average Jacobian values across time as a curve (referred to as **jCurves**). Thus, the shape changes of the heart are encoded by a similar number of curves as for eCurves, as LV are now represented by 16 jCurves and the RV by 11 jCurves.

The third and final type of encoding we use for comparison are the volumetric measurements frequently used in clinical settings, *i.e.*, we extract the volume (**Volume**) of each structure and time point directly from the corresponding segmentations masks.

3) Computing Accuracy Scores via Nested Cross-Validation—We infer the accuracy of a temporal shape encoding from agreement among the accuracy scores of three different classifiers in distinguishing healthy cases from diseased subjects based on that encoding. The three classifiers are a sparsity-constrained, logistic regressor (LOG-REG) [47] as well as Matlab's linear Support Vector Machine (LIN-SVM) [48] and non-linear Support Vector Machine with Gaussian Radial Basis Function kernel (RBF-SVM) [49]. We determine the optimal parameter setting and the testing accuracy of each classifier via nested two-fold cross-validation. To determine the optimal parameter setting, each classifier with a specific parameter setting is trained on 30 cases of the training fold and the remaining 14 cases of that fold are used to compute the training accuracy. The classifier with the highest training accuracy is then applied to the other fold in order to determine the testing accuracy. The process is repeated one more time after switching the training and testing fold. The remainder of this section will review in detail the parameter search space, which is defined with respect to the parameters specific of each classifier and temporal shape encoding, and the computation of the accuracy scores.

For LIN-SVM and RBF-SVM, we use the default parameters of their Matlab implementation (*e.g.*, the default implementation of RBF-SVM used the scaling factor $\sigma=1$) with the exception of the margin parameter C , whose search space is $\{-3, -1, \dots, 7\}$. We again use the default parameters for LOG-REG with the exception of the regularization parameter λ , whose search space is $\{0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1\}$.

With respect to the temporal shape encodings, the search space specific to jCurves is confined to the regularization parameter $s \in \{0.5, 1.5, 2.5\}$ associated with Demons based non-rigid registration [45]. For the eValue based measurements, the search space is defined with respect to the threshold $\tau \in \{50\%, 60\%, 70\%\}$. We determine the lower bound of this search space by carefully reviewing the tracking for the first 25 tracked modes across all time points and subjects, which, according to Figure 7, is the mode around which the robustness scores level off. The review reveals several example of pairs of eigenfunctions with inner products below 50% that show large differences (see examples in Figure 9). Thus, we set the lower bound of the search space to 50%.

Given τ , we determine the search space for N , the number of modes used by the classifier, based on the robustness score (*i.e.*, 5th percentile of the minimum inner product) with respect to the training data. Using the first fold for training, the search space for N is $\{4, 10, 18\}$ for the RV and $\{4, 12, 29\}$ for the LV. With respect to the second fold, the search space for N is $\{3, 7, 12\}$ for the RV and $N \in \{3, 19, 32\}$ for the LV.

After training classifiers, we compute their (training or testing) accuracy via the normalized accuracy score (nAcc):

$$\text{nAcc} \triangleq \frac{1}{2} \left(\frac{\text{\#of correctly labeled TOF case}}{\text{\#of TOF cases}} + \frac{\text{\#of correctly labeled healthy cases}}{\text{\#of healthy cases}} \right). \quad (11)$$

For simplicity, we refer to this score as *accuracy*. Note, that this accuracy score accounts for the asymmetry in the patient group sizes and is equivalent to the Area under the Curve for binary classifiers, such as SVM.

During testing, we also record the specificity and sensitivity of each classifier. In addition, we compute the Yates' Chi-Square test [50] and interpret the information stored by a temporal shape encoding significantly more informative than random noise if the corresponding p-values are below 2% for all three classifiers (LIN-SVM, RBF-SVM, and LOG-REG). For (target) encodings that are more informative than noise, we also record alternative encodings that are significantly less informative. Specifically, we confine the investigation to alternative encodings, which have lower nAcc scores for all three classifiers than the target encoding. We then compare those alternative encodings to the target encoding by computing the DeLong's Test [51] for each classifier type. If for all three classifiers the DeLong's Test return p-values below 2%, we view the target encoding significantly more informative than the alternative encoding.

B. Comparing Temporal Shape Encodings on the TOF Data Set

Table I summarizes the accuracy scores for the different temporal shape encodings of the left and right ventricles of 46 healthy controls and 42 TOF patients. The accuracy scores of all encodings agree with respect to being insignificantly better (gray) than chance for the LV and significantly better (black and red) than chance for the RV. These findings are also in agreement with the medical literature that reports that only the RV is impacted by TOF [52], [53], [54]. The remainder of this section focuses on the significant findings, *i.e.*, those of the RV.

For the RV, the accuracy scores for each encoding are consistent across the three classifiers (nAcc score vary by no more than 3.6%) with the exception of Volume. Volume not only achieves a high accuracy score (87.4%) with respect to the LIN-SVM but also is awarded the lowest accuracy score across all encodings and classifier for the RV, *i.e.*, the 76.8% with respect to the RBF-SVM. This instability in the performance of the classifiers indicates that the volumetric encoding contains, as expected, provides valuable information about the long-term impact of TOF but the information does not generalize well, *i.e.*, automatic inference based on the information is highly susceptible to the setting of the classifier. This

finding is not surprising when comparing the curves of the different encodings plotted in Figure 10. Volume is the only encoding characterizing the changes in the shape of the RV by a single curve, which limits the information that can be captured.

jCurves, on the other side, uses a set of curves to encode the shape changes. However, most of those curves highly fluctuate over time. This fluctuation does not reflect the subtle shape changes of the RV across time but rather indicates that the encoding is highly sensitive to noise. This sensitivity is also reflected in the accuracy scores of the three classifier whose average, unlike for the other three encodings, is below 80%. The fluctuations are most likely a result of the way the deformation maps of jCurves are computed. Log-Demons [45] requires finding the optimum to an objective functions composed of a data and regularization term. The minimum of this objective function needs to satisfy both criteria up to some extent. This trade-off decreases the sensitivity of the measures derived from the deformation fields. Defining an objective function just by a data term is a possible solution to this issue (see also the recently published registration methods based on shape signatures [55], [40]). Unfortunately, the popular implementation chosen for this experiment, *i.e.*, [45], only well-behaves in the presence of a regularization term.

An improvement over jCurves and Volume are the sensitivity and accuracy scores of eValues (at least 10.4% higher accuracy than jCurves and 3.8% higher than Volume). eValues also achieves at least 6.5% higher specificity than JCurves. The reason for the improvement can be seen in Figure 10, where eValues not only encodes temporal shape changes with multiple curves but these curves are also smooth as is the motion of the RV. The smoothness of the curves is not surprising since the eigenvalue sequence has explicit ties to the geometric invariant integrated over the surface of the object (see also Equation 3). By not sacrificing its representation power (for example via regularization), eValues yields a highly descriptive features set to distinguish the two cohorts. Note that the accuracy of eValues drops significantly when only applied to a single time point (not included in the table: nAcc=80.7% for LIN-SVM, nAcc=81.8% for RBF-SVM, nAcc=80.7% for LOG-REG) indicating that TOF not only impacts the shape of the RV but also its motion.

eCurves not only achieves the highest accuracy scores but is also the only encoding that always receives accuracy scores above 90%. It is also the only encoding that is significantly more accurate than another encoding, *i.e.*, jCurves, for all three classifiers ($p < 0.012$). Furthermore, Figure 10 shows that the tracking produces curves that are very different from those of eValues, especially those of higher order modes. Our results seem to indicate that these differences in the curves of eCurves lead to a more accurate temporal shape encoding as analyzing changes in the spectrum of the Laplacian operator across time is not only based on principled mathematical model but also results in a more accurately encoding of the shape changes due to TOF.

In summary, these experiments highlight the importance of eTracking and temporal consistency across measurements of a deforming shape. While there are a different ways of modeling these two issues, we achieve the highest accuracy for the RV by first computing the eValues for each time point and then performing eTracking across neighboring time points. Not only do eCurves record significant improvement over the deformation-based

encoding jCurves for the experiments involving the RV, but the technology needed for computing eCurves is also much simpler (a couple of lines of Matlab code). Finally, our results indicate that the long-term impact of reconstructive surgery of TOF not only is confined to the shape of the RV but also its motion, which may positively impact the decision process regarding follow up surgery for TOF patients.

IV. Future work

More refined discoveries would require improving the analysis, for example by adding additional shape measures to the analysis that can be explicitly linked to anatomical or functional changes. Alternatively, one could experiment with replacing the Laplacian spectrum with other popular shape encodings such as the 3D Zernike moments [56]. Furthermore, we could apply more sophisticated technology for matching the spectral signatures of the deforming heart chambers across time, such as those previously tested on matching 3D brain structures across subjects, whose shape can differ quite a lot [34], [35], [39], [36]. These technologies could also be applied for matching of curves across subjects to define the order of curves within a subject. Currently, this order is defined by the ascending ranking of eigenvalues at the first (end-diastole) time point. This increases the risk of biasing the analysis similar to the strategy in atlas construction of *a priori* selecting a fixed template [57]. Another possibility of reducing this risk is by tracking eigenpairs across time points taking both adjacent time points into account, *e.g.*, time point $i - 1$ and $i + 1$ for time point i . This problem relates to multi-way graph matching, which is considered a NP-hard problem whose solution can only be approximated. Finally, we could reduce the redundancy of eCurves, *i.e.*, making the representation more compact, by identifying characteristic eCurves across a data set via dimensionality reduction approaches (such as Principal Component Analysis) and then encode each individual curve of eCurves as a combination of characteristic curves. This would not only reduce the redundancy across the eCurves but also provide a mechanism for explicitly recording temporal change patterns as shown by our preliminary results in [37].

V. Conclusions

We present a new representation for temporal shape analysis that exploits the inherent properties of spectral signatures to capture morphological changes over time. We first compute the signatures (eValues) for each shape and at each time point. We then link the eValues between adjacent time points via eTracking, which is based on measuring the similarity of the eigenfunctions associated with the eValues. eTracking not only insures consistency of our shape scores across time points but also enables us to construct eCurves by following each eValue from the first to the last time point. The strength of our temporal shape lies in its simplicity. It only requires the time points within the series to be rigidly aligned to enable the tracking based on the spectral eigenfunctions. More importantly, it does not require prior registration across subjects since the spectral signature is pose invariant. In addition, eCurves only relies on one parameter: the number of modes for computing the spectrum, which we link to the lowest acceptable threshold for the inner product between eigenfunctions of linked modes. Finally, our representation does not involve statistical or physical modeling and only assumes a smooth transition of the shapes

between time points, so eCurves can potentially be used for a variety of applications in medical imaging.

We apply the temporal shape representation to encode the changes in the shape of the heart. We use eCurves to encode the shape changes of the right and left ventricles across 30 time points of 42 TOF patients and 46 healthy controls. We measure its sensitivity, specificity, and accuracy based on two-fold, nested cross-validation of three different classifiers and compare the classification results to those generated from series of 3D volumetric measurements and deformation-based encodings. All encodings receive insignificantly better accuracy scores than chance for the LV and significantly better ones for the RV, the structure most relevant for TOF. With respect to the RV, our proposed encoding achieves the highest accuracy scores and is significantly more accurate than the deformation-based encoding. Furthermore, our encoding provides new insight regarding the long term effects of TOF as it documents that the reconstructive surgery not only impacts the morphometry but also function of the right ventricle.

Acknowledgments

We would like to thank DongHye Ye for his help on generating the cardiac data-set. This project was supported in part by the NIH grants R01 HL127661, UL1 RR024134, and by the Institute for Translational Medicine and Therapeutics' (ITMAT) Transdisciplinary Program at Penn.

References

1. Morcos P, Vick I, Wesley G, Sahn D, Jerosch-Herold M, Shurman A, Sheehan F. Correlation of right ventricular ejection fraction and tricuspid annular plane systolic excursion in Tetralogy of Fallot by magnetic resonance imaging. *The International Journal of Cardiovascular Imaging*. 2009; 25(3): 263–270. [PubMed: 19048388]
2. Solarz D, Witt S, Glascock B, Jones F, Khoury P, Kimball T. Right ventricular strain rate and strain analysis in patients with repaired Tetralogy of Fallot: possible interventricular septal compensation. *Journal of the American Society of Echocardiography*. 2004; 17(4):338–44. [PubMed: 15044867]
3. Ances B, Ortega M, Vaida F, Heaps J, Paul R. Independent effects of HIV, aging, and HAART on brain volumetric measures. *Journal of Acquired Immune Deficiency Syndrome*. 2012; 59(5):469–77.
4. Pfefferbaum A, Rohlfing T, Rosenbloom MJ, Chu W, Colrain IM, Sullivan EV. Variation in longitudinal trajectories of regional brain volumes of healthy men and women (ages 10 to 85 years) measured with atlas-based parcellation of MRI. *NeuroImage*. 2013; 65:176–193. [PubMed: 23063452]
5. Ardekani S, Weiss RG, Lardo AC, George RT, Lima JAC, Wu KC, Miller MI, Winslow RL, Younes L. Cardiac motion analysis in ischemic and non-ischemic cardiomyopathy using parallel transport. *IEEE International Symposium on Biomedical Imaging*. 2009:899–902.
6. Beg MF, Helm PA, Mcveigh E, Miller MI, Winslow RL. Computational cardiac anatomy using MRI. *Magnetic Resonance in Medicine*. 2004; 52(5):1167–1174. [PubMed: 15508155]
7. Pearlman, PC., Igum, I., Kersbergen, KJ., Benders, MJNL., Viergever, MA., Pluim, JPW. Implicit surface registration with surface-oriented anisotropic deformation field smoothing. *IEEE International Symposium on Biomedical Imaging*; April 2013; p. 480–483.
8. Chetelat G, Landeau B, Eustache F, Mezenge F, Viader F, de la Sayette V, Desgranges B, Baron J-C. Using voxel-based morphometry to map the structural changes associated with rapid conversion in MCI: a longitudinal MRI study. *Neuroimage*. 2005; 27(4):934–46. [PubMed: 15979341]
9. Shen D, Davatzikos C. Measuring temporal morphological changes robustly in brain MR images via 4-dimensional template warping. *NeuroImage*. Apr; 2004 21(4):1508–1517. [PubMed: 15050575]

10. Thompson PM, Giedd JN, Woods RP, MacDonald D, Evans AC, Toga AW. Growth patterns in the developing brain detected by using continuum mechanical tensor maps. *Nature*. 2000; 404(6774): 190–3. [PubMed: 10724172]
11. Hoogendoorn C, Sukno F, Ordás S, Frangi A. Bilinear models for spatio-temporal point distribution analysis. *International Journal of Computer Vision*. 2009; 85(3):237–252.
12. Hoogendoorn C, Duchateau N, Sanchez-Quintana D, Whitmarsh T, Sukno F, De Craene M, Lekadir K, Frangi A. A high-resolution atlas and statistical model of the human heart from multislice CT. *IEEE Transactions on Medical Imaging*. Jan; 2013 32(1):28–44. [PubMed: 23204277]
13. Metz CM, Baka N, Kirisli H, Schaap M. Regression-based cardiac motion prediction from single-phase CTA. *IEEE Transactions on Medical Imaging*. Jun; 2012 31(6):1311–1325. [PubMed: 22438512]
14. Baloch S, Davatzikos C. Morphological appearance manifolds in computational anatomy: groupwise registration and morphological analysis. *Neuroimage*. Mar; 2009 45(1 Suppl):S73–S85. [PubMed: 19061962]
15. Raptis M, Soatto S. Tracklet descriptors for action modeling and video analysis. *European Conference on Computer Vision*. 2010:577–590.
16. Stauffer C, Grimson WEL. Learning patterns of activity using real-time tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2000; 22:747–757.
17. Matej P, Matej K, Kováčik S, Vučković G, Janež P. A trajectory-based analysis of coordinated team activity in a basketball game. *Computer Vision and Image Understanding*. 2009; 113(5):612–621.
18. Niebles JC, Chen C-W, Fei-Fei L. Modeling temporal structure of decomposable motion segments for activity classification. *European Conference on Computer Vision*. 2010:392–405.
19. Gorelick L, Blank M, Shechtman E, Irani M, Basri R. Actions as space-time shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. Dec; 2007 29(12):2247–2253. [PubMed: 17934233]
20. Batra D, Chen T, Sukthankar R. Space-time shapelets for action recognition. *IEEE Workshop on Motion and Video Computing*. 2008:1–6.
21. Ke Y, Sukthankar R, Hebert M. Spatio-temporal shape and flow correlation for action recognition. *IEEE Conference on Computer Vision and Pattern Recognition*. 2007:1–8.
22. Ma S, Zhang J, Ikizler-Cinbis N, Sclaroff S. Action recognition and localization by hierarchical space-time segments. *IEEE International Conference on Computer Vision*. 2013:2744–2751.
23. Dollar P, Rabaud V, Cottrell G, Belongie S. Behavior recognition via sparse spatio-temporal features. *Visual Surveillance and Performance Evaluation of Tracking and Surveillance*, 2005 2nd Joint IEEE International Workshop on; Oct 2005; p. 65-72.
24. Reuter M, Wolter F-E, Peinecke N. Laplace-Beltrami spectra as 'Shape-DNA' of surfaces and solids. *Computer-Aided Design*. 2006; 38:342–66.
25. Reuter M, Wolter F-E, Shenton M, Niethammer M. Laplacebeltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. *Computer-Aided Design*. 2009; 41:739–55. [PubMed: 20161035]
26. Weyl H. Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen. *Math Ann*. 1912:441–69.
27. Kac M. Can one hear the shape of a drum? *The American Mathematical Monthly*. 1966; 73(7):1–23.
28. Courant R, Hilbert D. *Method of Mathematical Physics*. Vol. I. Interscience Publishers; 1966.
29. Bronstein M, Bronstein A. Shape recognition with spectral distances. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2011; 33(5):1065–1071. [PubMed: 21135442]
30. Bronstein MM, Kokkinos I. Scale-invariant heat kernel signatures for non-rigid shape recognition. *IEEE Conference on Computer Vision and Pattern Recognition*. 2010:1704–1711.
31. Patane G, Spagnuolo M. Heat diffusion kernel and distance on surface meshes and point sets. *Computers & Graphics*. Mar; 2013 30(3):276–295.

32. Niethammer M, Reuter M, Wolter F-E, Peinecke N, Koo M-S, Shenton ME. Global medical shape analysis using the laplacebeltrami spectrum. *International Conference on Medical Image Computing and Computer Assisted Intervention*. 2007;850–857.
33. Konukoglu E, Glocker B, Criminisi A, Pohl KM. WESD – weighted spectral distance for measuring shape dissimilarity. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. Sep; 2013 35(9):2284–2297. [PubMed: 23868785]
34. Mateus, D., Horaud, R., Knossow, D., Cuzzolin, F., Boyer, E. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. *IEEE Conference on Computer Vision and Pattern Recognition*; June 2008; p. 1-8.
35. Shi Y, Lai R, Wang DJJ, Pelletier D, Mohr D, Sicotte N, Toga AW. Metric optimization for surface analysis in the laplacebeltrami embedding space. *IEEE Transactions on Medical Imaging*. Jul; 2014 33(7):1447–1463. [PubMed: 24686245]
36. Lombaert, H., Arcaro, M., Ayache, N. *Information Processing in Medical Imaging*. Springer International Publishing; 2015. Brain transfer: spectral analysis of cortical surfaces and functional maps; p. 474-487.
37. Bernardis, E., Konukoglu, E., Ou, Y., Metaxas, DN., Desjardins, B., Pohl, KM. Temporal shape analysis via the spectral signature. *International Conference on Medical Image Computing and Computer Assisted Intervention*; 2012; p. 49-56.
38. Baumgartner H, Bonhoeffer P, De Groot NM, de Haan F, Deanfield JE, Galie N, Gatzoulis MA, Gohlke-Baerwolf C, Kaemmerer H, Kilner P, Meijboom F, Mulder BJ, Oechslin E, Oliver JM, Serraf A, Szatmari A, Thaulow E, Vouhe PR, Walma E, Vahanian A, Auricchio A, Bax J, Ceconi C, Dean V, Filippatos G, Funck-Brentano C, Hobbs R, Kearney P, McDonagh T, Popescu BA, Reiner Z, Sechtem U, Sirnes PA, Tendera M, Vardas P, Widimsky P, McDonagh T, Swan L, Andreotti F, Beghetti M, Borggrefe M, Bozio A, Brecker S, Budts W, Hess J, Hirsch R, Jondeau G, Kokkonen J, Kozelj M, Kucukoglu S, Laan M, Lionis C, Metreveli I, Moons P, Pieper PG, Pilosoff V, Popelova J, Price S, Roos-Hesselink J, Uva MS, Tornos P, Trindade PT, Ukkonen H, Walker H, Webb GD, Westby J. Esc guidelines for the management of grown-up congenital heart disease (new version 2010): The task force on the management of grown-up congenital heart disease of the european society of cardiology (esc). *European Heart Journal*. 2010; 31(23):2915. [Online]. Available: [+http://dx.doi.org/10.1093/eurheartj/ehq249](http://dx.doi.org/10.1093/eurheartj/ehq249). [PubMed: 20801927]
39. Ovsjanikov M, Ben-Chen M, Solomon J, Butscher A, Guibas L. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics*. 2012; 31(4):30.
40. Lombaert H, Grady L, Polimeni JR, Cheriet F. Focusr: Feature oriented correspondence using spectral regularization-a method for precise surface matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2013; 35(9):2143–2160. [PubMed: 23868776]
41. Protter M. Can one hear the shape of a drum? Revisited. *SIAM Review*. Jun; 1987 29(2):185–197.
42. Ye D, Desjardins B, Hamm J, Litt H, Pohl K. Regional manifold learning for disease classification. *Medical Imaging, IEEE Transactions on*. Jun; 2014 33(6):1236–1247.
43. Ou Y, Sotiras A, Paragios N, Davatzikos C. Dramms: Deformable registration via attribute matching and mutual-saliency weighting. *Medical Image Analysis*. 2011; 15(4):622–639. [PubMed: 20688559]
44. Kleinberg, J., Tardos, E. *Algorithm Design*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc; 2005.
45. Vercauteren T, Pennec X, Perchant A, Ayache N. Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage*. 2009; 45(1):S61–S72. [PubMed: 19041946]
46. Cerqueira MD, Weissman N, Dilsizian V, Jacobs A, Kaul S, Laskey W, Pennell D, Rumberger J, Ryan T, Verani M. Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart. a tatement for healthcare professionals from the cardiac imaging committee of the council on clinical cardiology of the american heart association. *Circulation*. 2002; 105:539–542. [PubMed: 11815441]
47. Liu, J., Ji, S., Ye, J. SLEP: Sparse Learning with Efficient Projections, Arizona State University. 2009. [Online]. Available: <http://www.public.asu.edu/~jye02/Software/SLEP>
48. Burges CJC. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*. 1998; 2:121–167.

49. Cortes C, Vapnik V. Support-vector networks. *Machine Learning*. 1995; 20(3):273–297. [Online]. Available: <http://dx.doi.org/10.1007/BF00994018>.
50. Yates F. Contingency tables involving small numbers and the χ^2 test. Supplement to the *Journal of the Royal Statistical Society*. 1934; 1(2):217–235. [Online]. Available: <http://www.jstor.org/stable/2983604>.
51. DeLong E, DeLong D, Clarke-Pearson D. Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*. Sep; 1988 44(3): 837–845. [PubMed: 3203132]
52. Norton KI, Tong C, Glass RBJ, Nielsen JC. Cardiac MR imaging assessment following Tetralogy of Fallot repair. *Radio Graphics*. 2006; 26:197–212.
53. Zhang H, Wahle A, Johnson RK, Scholz TD, Sonka M. 4-D cardiac MR image analysis: Left and right ventricular morphology and function. *IEEE Transactions on Medical Imaging*. 2010; 29(2): 350–364. [PubMed: 19709962]
54. Bodhey NK, Beerbaum P, Sarikouch S, Kropf S, Lange P, Berger F, Anderson RH, Kuehne T. Functional analysis of the components of the right ventricle in the setting of Tetralogy of Fallot. *Circular Cardiovasc Imaging*. Sep; 2008 1(2):141–7.
55. Lombaert H, Grady L, Pennec X, Ayache N, Cheriet F. Spectral log-demons: Diffeomorphic image registration with very large deformations. *International Journal of Computer Vision*. 2014; 107(3): 254–271. [Online]. Available: <http://dx.doi.org/10.1007/s11263-013-0681-5>.
56. Novotni, M., Klein, R. Proceedings of the Eighth ACM Symposium on Solid Modeling and Applications. New York, NY, USA: ACM; 2003. 3D Zernike descriptors for content based shape retrieval; p. 216-225. ser SM '03[Online]. Available: <http://doi.acm.org/10.1145/781606.781639>
57. Zöllei L, Shenton M, Wells W, Pohl K. The impact of atlas formation methods on atlas-guided brain segmentation, statistical registration. Pair-wise and Group-wise Alignment and Atlas Formation Workshop at MICCAI 2007: Medical Image Computing and Computer-Assisted Intervention. 2007:39–46.




Fig. 1. 3D models representing the left (green) and right ventricle (red) at different time points in the cardiac cycle according to a single cine MR scan of a healthy patient. Notice the relative small changes in the shape of the anatomy between adjacent time points.

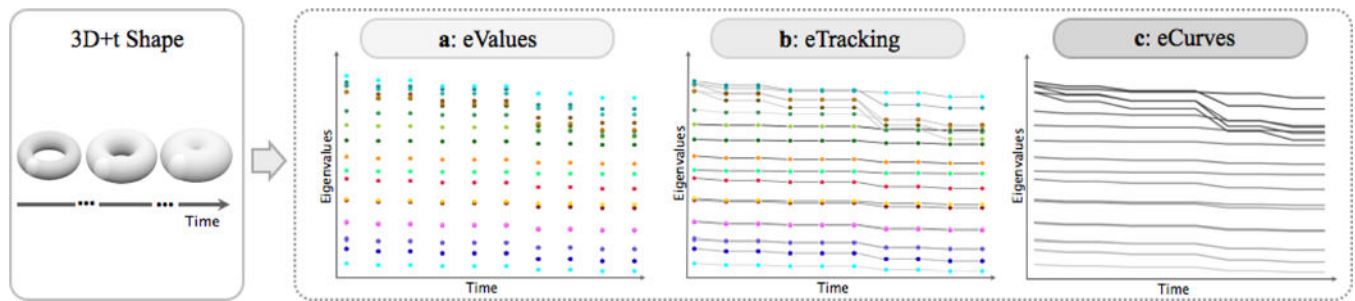
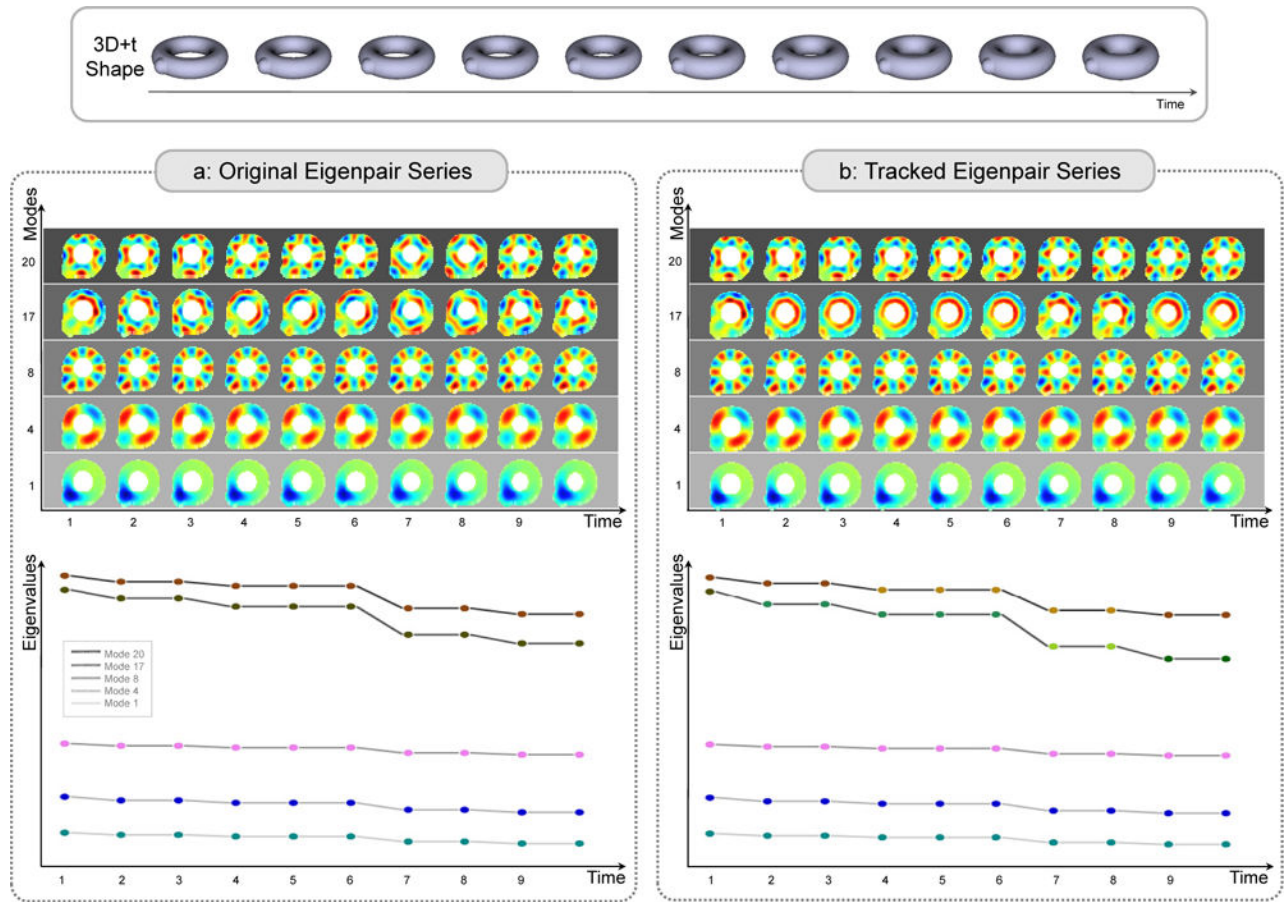


Fig. 2.

Encoding a series of 3D shapes overview. **a:** First, we encode the 3D shape at each time point by computing the eigensystem of the Laplace operator applied to the corresponding segmentation (eValues). Each column corresponds to the eigenvalues of one time point in the image sequence. Colors correspond to the modes at which the eigenvalue was computed, *i.e.*, the eigenvalues of n^{th} mode are plotted in the same color across time. While with respect to a single time point the order of the eigenvalues is clearly defined, this order is not guaranteed to hold across time point, *e.g.*, the n^{th} mode at time t does not necessarily relate to the n^{th} mode at time $t + 1$. **b:** We track eigenvalues between adjacent time points by pairing those whose eigenfunctions are most similar (eTracking). Tracking of the modes across time points is delineated with gray lines, the darker the line, the higher the mode. Original mode colors are left to highlight the effects of tracking. To focus on tracking, eigenvalues show in 'a:' are omitted in 'b:' if they were not part of the first 16 tracked modes. **c:** Finally, following the tracked eigenvalues, we obtain eigenvalue Curves (eCurves). eCurves represent the shapes and its changes of the object in a.

**Fig. 3.**

4D eigenseries of a small bump attached to a torus, whose radius of the tube is increasing with time (*i.e.*, the 10 time points). **a:** Eigenvalues $\Lambda(t)$ and eigenfunctions $F(t)$ for increasing eigenmodes $j = 1, 4, 8, 17, 20$. **b:** The same modes after eTracking. While the lower modes (1,4,8) are equivalent to a), the eigenfunctions associated with mode 17 and 20 now smoothly vary across time, properly reflecting the deforming shape. Interestingly, the corresponding eigenvalues now change more drastically over time compared to the original series of eigenvalues.

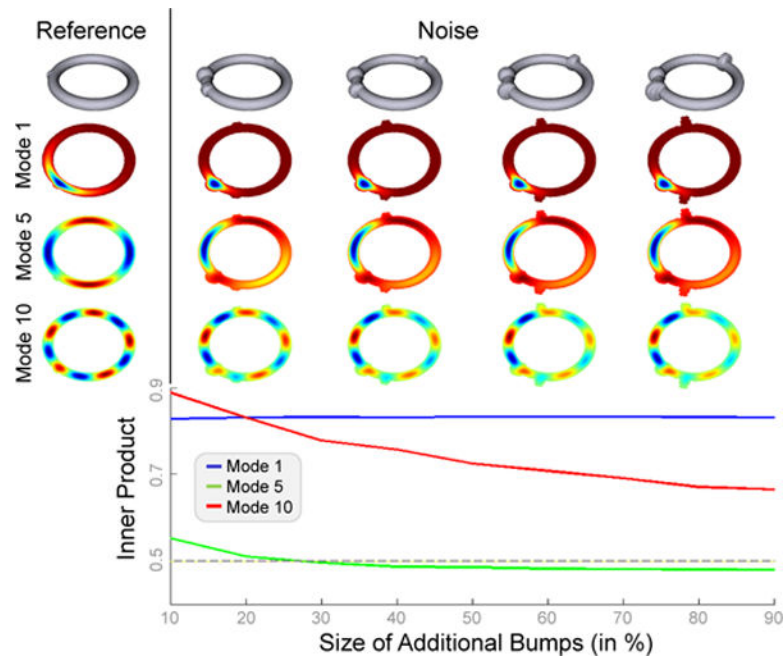


Fig. 4.

The impact of small shape changes on the inner product between the corresponding eigenfunctions. The shape on the left (reference) is a torus with an additional small bump (bump of interest). The tori to the right have a larger bump at that location simulating a small change of the original shape. To model a source of noise, the torus includes two additional, smaller bumps, whose height increases from 10% to 90% with respect to the radius of the torus. For all these shapes, the eigenfunctions of the first, fifth and tenth mode are shown below. Their inner products to the corresponding eigenfunctions of the reference (on the left) are plotted in the graph below. Note, the eigenfunctions of higher modes (*e.g.*, Mode 10) are more sensitive towards small shape changes than those of lower modes (*e.g.*, Mode 1). However, the eigenfunctions of the lower modes (*e.g.*, Mode 1) are also impacted by small changes, which is difficult to visualize due to the limiting color spectrum but is implied by the corresponding plot of the inner product being slightly slanted.

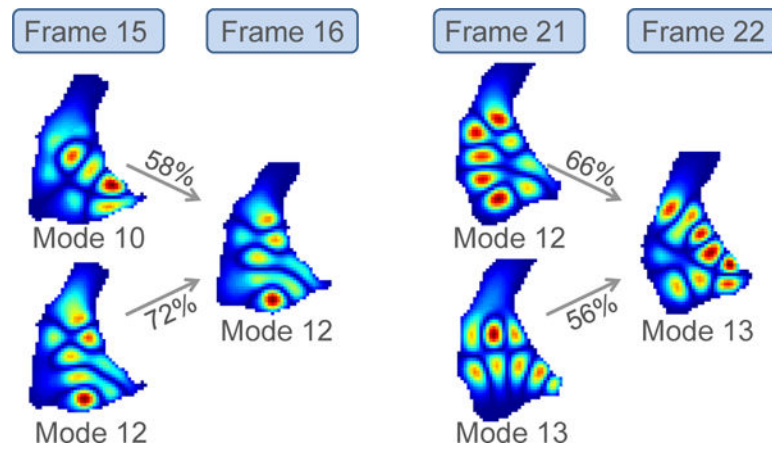


Fig. 5.

Examples of two modes that have their highest inner product (in percent on the left: 58% and 72%; in percent on the right: 66% and 56%) across all modes at the next time point with the same mode. In the example on the left, mode 12 at time point 15 will be linked to mode 12 at time point 16, while in the right example mode 12 at time point 21 will be linked to mode 13 at time point 22.

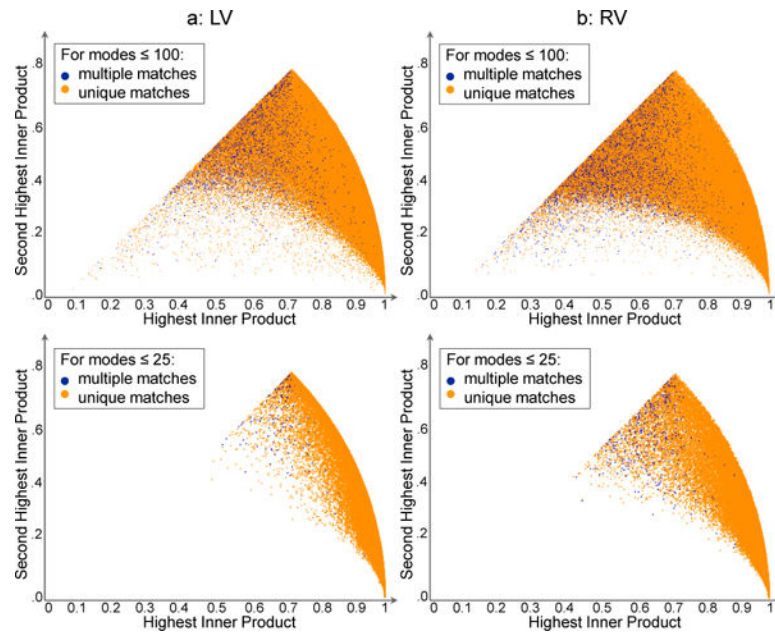


Fig. 6.

Plotting modes according to highest and second highest inner product with modes of the previous time points for **a: LV** and **b: RV**. Uniquely matched modes are marked in orange while in blue are multiple matched mode, *i.e.*, the ones for which two or more modes of the previous time point have a highest inner product with the same mode at the next time point. The first row plots all 100 modes across all data sets. In all cases, the number of multiple matches is less than 7% when considering all modes. The second row shows the same plot for the first 25 modes of each subject. The percentage of multiple matched modes now decreases by up to a factor of 10.




Fig. 7.

For the first 100 modes, the plot shows the robustness measure of the tracking, *i.e.*, the 5th percentile of the minimum inner product across all subjects. The score levels of around 40% at the 25th mode and then drops off around the 90th mode. Then setting the threshold τ to 50%, the optimal N would be 15 based on all samples.

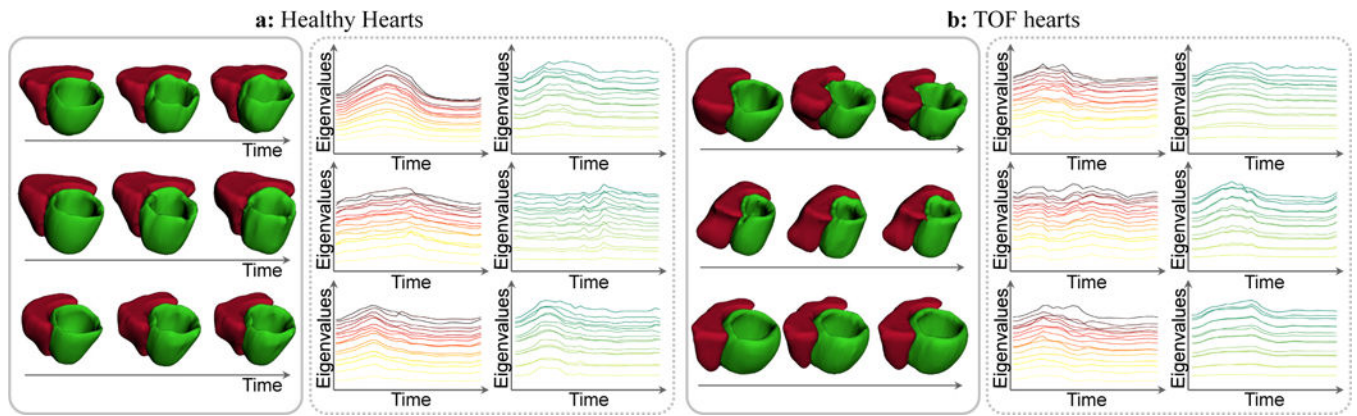


Fig. 8.

The 3D models of the RV (red) and LV (green) at three different time points of three healthy and three TOF subjects. The corresponding eCurves of the first 15 modes are plotted to the right. The eCurves in the red color scheme encode the series of RV shapes and the ones in the green color scheme encode the series of LV shapes. The end diastole corresponds to the first time point, while the peak in the eCurves corresponds to the end systole of the cardiac cycle. In comparison to the 3 healthy hearts, the peaks in the eCurves of the 3 TOF hearts seem less prominent in the right ventricle.

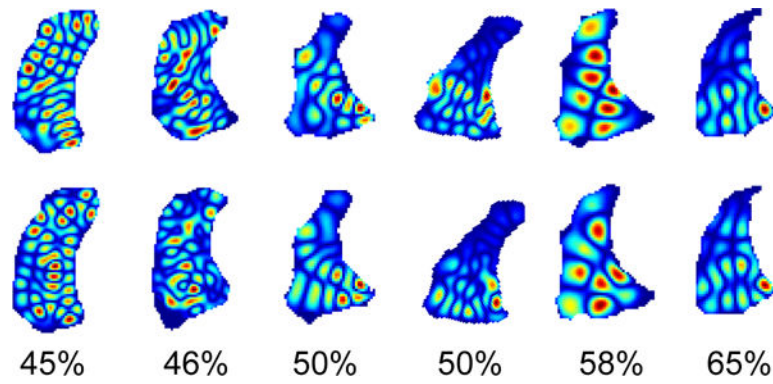


Fig. 9.

Eigenfunctions of the RV between two matched modes (first and second row) and associated inner product (third row) in six different cases. Note the large discrepancy in the patterns defined by the eigenfunctions below 50%, while their differences decrease as their inner product increases.




Fig. 10.

Four different encodings capturing the temporal changes in the shape of the RV in one case. While Volume confines capturing the shape changes by single, smooth curve, jCurves potentially can capture more complex shape changes by relying on multiple curves. However, some of those curves highly fluctuate over time, which is not in agreement with the rather small changes observed in the RV shape between neighboring time points. The only representations that capture shape changes through multiple smooth curves are those based on the signatures, *i.e.*, eValues and eCurves, which also explains the higher accuracy scores of those encodings. The consistency across time (tracking) enforced by eCurves results in curves that deviate from those of eValues (especially those associated with higher modes). These deviations seem to reflect more accurately the temporal shape changes of the RV, as eCurves achieves the overall highest accuracy scores.

TABLE I

Specificity, sensitivity, and normalized accuracy scores (nAcc) of the encodings in % with respect to a linear SVM (LIN-SVM), non-linear SVM (RBF-SVM), and logistic regression (LOG-REG) classifier based on **a**: volumetric measurements; **b**: jCurves; **c**: Eigenvalues directly (eValues); **d**: eCurves with tracking. Results are marked in light gray if they were not and black if they were significantly better than chance for all three classifiers. Results in red are significantly better than jcurves for all three classifiers.

	LIN-SVM						RBF-SVM						LOG-REG					
	Sen	Spec	nAcc	Sen	Spec	nAcc	Sen	Spec	nAcc	Sen	Spec	nAcc	Sen	Spec	nAcc	Sen	Spec	nAcc
a: Volume				85.7		87.4				66.7		76.8				88.1	82.6	85.4
b: jCurves				83.3		80.8				76.2		78.3				76.2	78.3	77.2
c: eValues				97.6		91.2				92.9		88.8				90.5	89.1	89.9
d: eCurves																		