



HAL
open science

Statistical Shape Analysis of Surfaces in Medical Images Applied to the Tetralogy of Fallot Heart

Kristin Mcleod, Tommaso Mansi, Maxime Sermesant, Giacomo Pongiglione,
Xavier Pennec

► **To cite this version:**

Kristin Mcleod, Tommaso Mansi, Maxime Sermesant, Giacomo Pongiglione, Xavier Pennec. Statistical Shape Analysis of Surfaces in Medical Images Applied to the Tetralogy of Fallot Heart. Fred-eric Cazals and Pierre Kornprobst. Modeling in Computational Biology and Biomedicine, Springer, pp.165-191, 2013, Lectures Notes in Mathematical and Computational Biology, 978-3-642-31207-6. 10.1007/978-3-642-31208-3_5. hal-00813850

HAL Id: hal-00813850

<https://inria.hal.science/hal-00813850>

Submitted on 1 Jul 2013

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

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

Chapter 5

Statistical Shape Analysis of Surfaces in Medical Images Applied to the Tetralogy of Fallot Heart

Kristin McLeod, Tommaso Mansi, Maxime Sermesant, Giacomo Pongiglione, and Xavier Pennec

5.1 Introduction

During the past ten years, biophysical modeling of the human body has been a topic of increasing interest in the field of biomedical image analysis. The aim of such modeling is to formulate personalized medicine where a digital model of an organ can be adjusted to a patient from clinical data. This virtual organ would enable to estimate parameters which are difficult to quantify in clinical routine and to test therapies in-silico as well as to predict the evolution of the organ over time and with therapy.

While biophysical modeling has been successfully developed and validated at the cell level, the integration at the organ level is still challenging because the interaction between the physiology and the organ shape is complex. There is thus a strong need to understand the organ shape characteristics and their evolution at the population level, which is the topic of Computational Anatomy. By shape, we intend here, and in the sequel of this paper, geometric data such as curves and surfaces (interfaces of organs) that describe the anatomy. On top of this, Computational Physiology is aiming at describing and modeling the function of these organs and how they interact together. Anatomy can be observed in most medical image modalities, such as Computational Tomography (CT) or Magnetic Resonance Imaging (MRI), while measuring the physiology requires more specific imaging modalities that are often more difficult to acquire in a clinical context. The expectation is that anatomical shapes can then be used to better understand physiology, or at least some of its pathologies.

The biological mechanisms involved in the formation and evolution of the shape of organs are often too complex to model. Thus, one has to rely on robust statistical approaches to produce generative models from data. Given the shapes of a number of subjects, we can create a generalized model of an organ, called an atlas, which can be used to distinguish normal from pathological organs or can be deformed to give patient-specific models driven by the clinically available parameters of interest. Such statistical analyses can both provide a predictive model and guide the biophysical approach. However, computing statistics on dynamic 3D shapes is very challenging. Traditional methods rely on point based parameterizations of the shapes where the point-to-point correspondences can be an

Kristin McLeod

INRIA Sophia Antipolis Méditerranée, Asclepios project-team, 2004 Route des Lucioles, 06902 Sophia Antipolis Cedex, France, e-mail: Kristin.Mcleod@inria.fr

Xavier Pennec

INRIA Sophia Antipolis Méditerranée, Asclepios project-team, 2004 Route des Lucioles, 06902 Sophia Antipolis Cedex, France, e-mail: Xavier.Pennec@inria.fr

Tommaso Mansi

Siemens Corporate Research, Image Analytics and Informatics, Princeton, NJ, U.S.A, e-mail: Tommaso.Mansi@tmansi.net

Maxime Sermesant

INRIA Sophia Antipolis Méditerranée, Asclepios project-team, 2004 Route des Lucioles, 06902 Sophia Antipolis Cedex, France, e-mail: Maxime.Sermesant@inria.fr

Giacomo Pongiglione

Ospedale Pediatrico Bambine Gesù, Rome, Italy,

important limiting factor for the usability of the method since the correspondences must be correct and consistent over all the shapes. New approaches were recently developed to compute such statistics without this limitation [25, 26, 24, 6, 32, 15, 33, 10]. Here, the statistical shape analysis tools are based on currents, a non-parametric representation of shapes. These tools have a wide range of applications, and provide a well-posed framework for statistical shape analysis of groups. Due to the fact that the methods do not assume point correspondences between structures (and in fact assigning landmarks to structures such as the heart are arbitrary), a wider range of data can be used. For example, one can use surfaces to model organs such as the heart, brain, and lungs, curves to model sulcal lines on the brain cortex, and sets of curves to represent fibre bundles from diffusion MRI in the brain. The goal of this chapter is to give an overview of this methodology in the context of a specific clinical problem: the prediction of the cardiac shape remodeling in repaired tetralogy of Fallot patients due to chronic regurgitation. The underlying assumption is that changes in heart morphology may reveal structural abnormalities and dysfunctions due to the pathology. In particular, we aim to establish which shape patterns are related to the pathology in order to give further insights into the condition.

5.1.1 Repaired Tetralogy of Fallot

Tetralogy of Fallot (ToF) is a congenital heart defect that affects approximately four out of every 10,000 babies [9]. The primary defect associated with this condition is a ventricular-septal opening which allows blood to flow freely between the ventricles. Secondary is stenosis in the pulmonary artery which restricts the blood flow from the right ventricle to the lungs. Due to a misalignment of the aorta over the ventricular-septal defect, the aorta is fed by both the left and right ventricle rather than just the left ventricle. Patients may also have hypertrophy in the right ventricle that causes a boot-like shape of the ventricle which is characteristic of this condition.

These abnormalities require open heart surgical repair early in infancy. As part of this surgery, the stenosis in the pulmonary artery is cleared to allow blood to flow more freely through the artery. As a consequence, the pulmonary valves which control one way blood flow from the right ventricle to the pulmonary artery can be destroyed completely or damaged, causing blood to leak back to the right ventricle. The long-term outcome of this condition is right ventricular re-modeling caused by the regurgitated blood (Fig. 5.1). Nowadays, patients undergo follow up operations to reduce this post-operative sequelae. In particular, the valves are replaced with an artificial device to reduce the regurgitation. A clinical challenge for these patients is determining the optimal time for intervention. On the one hand, it is preferable to wait as long as possible before performing such operations since the artificial devices have a limited life-span resulting in the need for repeat follow-up operations. Either open heart surgery or percutaneous valve replacement may be required, depending on the diameter of the pulmonary annulus. These therapies carry some danger, so minimizing the frequency and total number of surgeries is crucial for these patients. On the other hand, the heart undergoes re-modeling due to the regurgitated blood producing volume overload in the right ventricle and if left long enough this damage can be irreversible. Determining the optimal time for intervention is a trade-off between maximizing the time between operations and minimizing damage to the heart.

Therefore, the clinical application for this condition is to better understand how the heart evolves over a large period of time, often decades, with and without surgery, with the main goal to be able to predict the optimal time for intervention. However, due to the large variability of shape, the optimal time, placement and size of the artificial device can be difficult to determine. Furthermore, the complexity of the biological mechanisms involved in heart growth hinders the development of a direct model of cardiac re-modeling. Statistical shape analysis can therefore be employed to aid in further understanding the pathology to assist cardiologists with diagnosis, therapy planning, and long term prognosis. More particularly, we are interested in determining quantitative measures of the shape that correlate with the cardiac function and the severity of the disease in these patients (bio-markers of the disease). Such bio-markers could be used to determine the severity of the disease when regurgitation cannot be assessed or to provide structural information that may appear before deterioration of the blood dynamics.




Fig. 5.1 Comparison between (a) a normal heart and (b) a Tetralogy of Fallot heart.

5.1.2 Chapter Overview

Public

This chapter is targeting people interested in statistics of surfaces, the surface being considered as the random variable itself. The reader is assumed to have good notions of multivariate statistics (expectation, moments, PCA, regression, hypothesis tests) and some notions of analysis in infinite dimensional vector spaces, in particular with reproducing kernel Hilbert spaces (RKHS). This last notion is borrowed from machine learning theory and is applied here to the characterization of geometric shapes through geometric integration theory. A few notions from Riemannian geometry are used but no specific knowledge is required.

Outline

In the following sections, we will first review in Sect. 5.2 statistical shape analysis methods with a particular focus on the statistical analysis of surfaces. We will introduce the formalism of deformations which is at the center of most of the current works in computational anatomy [19]. Then, we will detail the framework of currents to represent surfaces, and show how this can be turned into an effective shape analysis technique. Section 5.3 will apply this methodology to the shape of the heart of 13 patients with repaired tetralogy of Fallot. Correlating shape with clinical variables will illustrate how we can extract some insight about the relationship between morphology and physiology. Last but not least, we will exemplify how the lack of longitudinal measurements can be bypassed by building a statistical generative growth model from cross-sectional data which summarizes the heart shape remodeling at the population level.

5.2 Statistical Shape Analysis

5.2.1 Shapes, Forms and Deformations

There is generally no physical model that can faithfully relate the shape of organs in different patients. Thus, to analyze their variability in a population, one usually extracts some anatomically representative landmarks (or more generally geometric features), and models their statistical distribution across the population, via a mean shape and covariance structure analysis after a group-wise matching for

instance. One of the earliest methods [2, 3] consists in studying the variability of anatomical landmark positions among a population: after a global pose (position and orientation) normalization and the consistent identification of landmarks within all patients (Fig. 5.2), a principle component analysis (PCA) is performed to extract the main modes of variation of the shape.




Fig. 5.2 Surface correspondences. Using landmark-based methods to register one surface mesh to another surface mesh, for each point x in patient A , we require the corresponding point $y(x)$ in patient B . This requires a full parameterization of the patient surface meshes. This is illustrated on the right ventricle of two patients with rToF.

5.2.1.1 Shapes

In such a process, the global pose of our objects is often considered as a nuisance factor related to the arbitrary coordinate system of the acquisition device: generally speaking, the shape of an object is understood as the geometric information that is invariant under translation, rotations and rescaling. Thus, one is often interested only in the *shape*, i.e. what remains if we quotient out the object space by a group action, usually rigid body, similarity or affine transformations. For instance, when we consider the equivalent classes of sets of N (labeled or unlabeled) points under the action of a global similarity (resp. rigid body) transformation, we obtain the celebrated shape (resp. size and shape) spaces of Kendall et al [12, 4, 13, 31]. One can define similarly the shape spaces of curves by removing the effect of re-parameterizations and the global pose of the curve [27, 38, 29, 28]. However, in medical image analysis, the natural coordinate system of an organ is its position and orientation with respect to surrounding tissues and organs in the image. Thus, once images are normalized by aligning them to a reference anatomy, the pose of organs is a variable of interest and should be kept for analysis.

5.2.1.2 Deformations

An alternative modeling of shapes was proposed by D’Arcy Thompson in 1917 [35]. The idea is to assume that there is a template object, which represents the reference shape. In medical image analysis, the template is most often built from real observations and is called an atlas¹. The atlas can be seen as an equivalent of the mean shape. Then, the variability of the shape is analyzed through the deformations of this reference object towards the actual observations: a shape difference is encoded by the transformation that deforms one onto the other. As the deformation of a smooth object should

¹ In this paper, atlas is always taken in the sense of template and not in the sense of the atlas of differential geometry.

be a smooth object, we have to work with diffeomorphisms (invertible, one-to-one mappings with smooth inverses). This formalism was promoted to a generic shape analysis tool by Grenander and Miller [8, 18] based on advanced mathematical tools to compute on infinite dimensional groups of diffeomorphisms [36]. One key feature of this lift of the shape characteristics from the object space to the transformation space is that it allows to apply the typical deformations to other objects than the ones analyzed in the first place (provided that they also live in the coordinate system of the atlas). For instance, the surface of the cortex in the brain has deep foldings called *sulci* which are important anatomical features. [26] analyzed the variability of the trace of these sulci on the surface of the brain cortex to extract the main 3D deformation modes. These modes could then be used to deform accordingly the surface of the cortex or the full 3D volume of the brain.




Fig. 5.3 Riemannian statistics on the sphere: the Fréchet mean is the point minimizing the squared Riemannian distance. It corresponds to the point for which the development of the geodesics to the data points on the tangent space is optimally centered (the mean in that tangent space is zero). The covariance matrix is then defined in that tangent space. Figure adapted from [20].

However, the problem is even more complex than previously as we want here to perform statistics on large deformations, which are known to belong to a non-linear and infinite dimensional manifold. The Riemannian setting is one of the most powerful structures to generalize simple statistics to non-linear spaces: it provides a definition of the distance between points of our manifold and a notion of shortest path / straight lines using geodesics [20, 21]. The main difficulty is that the mean value cannot be defined through an integral or a sum as in Euclidean spaces. Instead, one must look for points in the manifold that minimize the dispersion of the other points around it, which is commonly measured by the variance (the mean squared distance). This is what is called the Fréchet or Karcher mean (Fig. 5.3).

Then, one can compute the covariance matrix (the directional dispersion around the mean) by developing the manifold onto its tangent space at the mean point. Basically, this amounts to representing each data point by the momentum (initial speed vector) needed to shoot² a geodesic to it from the mean.

This generalization of statistics has to be slightly modified to fit the atlas deformation model: the ‘distance’ between the atlas and a shape is given here by the length of the shortest path in the space of deformations. The atlas is called unbiased if it is centered, i.e. if it minimizes (among all possible atlases) the sum of squared distances to the shapes [30]. In this sense, this is the Fréchet mean of the shapes. Then, each shape is represented by the momentum of the deformation that allows to regenerate the shape by deforming the atlas. These momentum all belong to the tangent space³ of the diffeomorphisms at the identity, which is a vector space in which we can perform a PCA [14]: shooting in the space of deformations along the first eigenvectors of the covariance matrix of the momentum gives deformation modes that represent the main shape variability when applied to the atlas.

² We use the term “geodesic shooting” to define the integration of the Euler-Lagrange equations, which plays the role of the exponential map in Riemannian geometry

³ More precisely to the cotangent space

5.2.2 From Points to Surfaces: the Formalism of Currents

Landmarks can be encoded by the probability of their location in space (typically a Gaussian around their expected value). When the noise is going to zero, then this probability density function (pdf) becomes singular. However, we can continue to deal with these types of singularities by considering distributions (generalized functions) instead of functions, which include Diracs. Mathematically, a Dirac is not a function, but an object which can be characterized by the result of its integration against any function of a sufficiently smooth functional space: $\forall f \in W, \int \delta_x(y).f(y).dy = f(x)$. In a finite dimensional Euclidean vector space, any linear form ϕ (function from the vector space to \mathbb{R}) can be characterized using the scalar product: Thanks to the Rietz representation theorem, there exists a vector x such that $\phi(w) = \langle x | w \rangle$. The equivalent in infinite dimensional functional spaces is the Dirac distribution, which is actually an element of the linear functionals over the space W . In that framework, a set of N points x_k can be represented by the ‘pdf’ $p(x) = 1/N \sum_i \delta_{x_i}(x)$, and its evaluation on a function $f(x)$ results in the mean of the values of f at the points x_k .

When we move to curves and surfaces, we are dealing with the locally singular locus of points in certain directions only: curves and surfaces are continuous along their tangent. The extension of the notion of distributions that allows to take that information into account is the geometric integration theory. The basic idea is to integrate differential forms, but for objects like smooth surfaces in 3D, we shall simply define currents by their action on vector fields, similarly to the way distributions are defined by their action on scalar functions. In fact, one could see currents as the characterization of surfaces by the flux of vector fields through it.

5.2.2.1 Surfaces seen as Currents

Let W be a Hilbert space of test vector fields. This is a possibly infinite complete vector space provided with a scalar product. For a given surface S , we can measure the flux of any vector field $\omega \in W$ through this surface:

$$\mathcal{S}(\omega) = \int_{x \in S} \langle \omega(x) | n(x) \rangle d\sigma(x),$$

where $n(x)$ is the normal to the smooth surface at point $x \in S$ and $d\sigma(x)$ the surface element around point x . The shape of the surface S is characterized by the variation of the flux as the test vector field ω varies in W . Thus, the surface actually defines a continuous linear form on W which can be identified as an element of the dual space W^* , which is the space of linear functionals from W to \mathbb{R} (currents). The nice property of currents is that it is a vector space: we can add or subtract currents from each others, or multiply them by a scalar. However, we should keep in mind that the space of currents is larger than the space of smooth surfaces: one can for instance add many pieces of surfaces together in a non-continuous way to create a non-continuous object.

5.2.2.2 The Kernel Trick for Currents

Now that we have identified surfaces to currents, we need to define more carefully what is the space W that we consider. The core element proposed by [7] is to consider a kernel metric (typically the Gaussian kernel $K_W(x, z) = \exp(-\|x - z\|^2/\lambda_W^2)$) in order to turn W into a Reproducible Kernel Hilbert Space (RKHS). The reproducibility property implies that W is the dense span of basis vector fields of the form $\omega_x^\alpha(\cdot) = K_W(x, \cdot)\alpha$. This means that any vector field of W can be written as an infinite linear combination of such basis vectors. The kernel induces a scalar product which is easily computed on two basis vectors: $\langle \omega_x^\alpha | \omega_y^\beta \rangle_W = \alpha^T K_W(x, y)\beta$.

Since W is a Hilbert space, by the Rietz representation theorem, there exists an isometric linear mapping \mathcal{L}_W from W to W^* which maps to each $\omega \in W$ the linear form that reproduces the scalar product: $(\mathcal{L}_W(\omega))(\omega') = \langle \omega | \omega' \rangle_W$. This canonical isomorphism allows to define the dual of an element, or conversely to map to each surface the vector field that optimizes the flux. For instance, the dual of the above basis vectors are the Dirac delta currents $\delta_x^\alpha = \mathcal{L}_W(\omega_x^\alpha)$ as we have $\langle \omega_x^\alpha | \omega \rangle_W = \alpha^T \omega(x)$. These can be seen as vector fields whose spatial support is concentrated at one point only.

The space of currents W^* is the dense span of these basis elements. For instance, the surface S is represented by the current $\mathcal{S} = \int_{x \in S} \delta_x^{n(x)} d\sigma(x)$. The usual way to define a norm on the dual space is to take the operator norm: $\|\mathcal{S}\|_{W^*} = \sup_{\omega \in W, \|\omega\|=1} |\mathcal{S}(\omega)|$. The distance induced by this norm might seem difficult to use, but thanks to the RKHS properties, we have a closed form for basis vectors:

$$\langle \delta_x^\alpha \mid \delta_y^\beta \rangle_{W^*} = \alpha^\top K_W(x, y) \beta \quad \text{and} \quad \|\delta_x^\alpha - \delta_y^\beta\|_{W^*}^2 = \|\delta_x^\alpha\|_W^2 + \|\delta_y^\beta\|_W^2 - 2 \langle \delta_x^\alpha \mid \delta_y^\beta \rangle_W.$$

5.2.2.3 Numerical Issues

It is interesting to notice that this distance can be approximated by

$$\|\delta_x^\alpha - \delta_y^\beta\|_{W^*}^2 = \|\alpha - \beta\|^2 + 2\|x - y\|^2 / \lambda_W^2 \alpha^\top \beta + O(\|x - y\|^4 / \lambda_W^4),$$

when the points x and y are within a fraction of λ_W , while the distance is essentially constant (and equal to $\|\alpha\|_{L_2}^2 + \|\beta\|_{L_2}^2$) when the points x and y are more than a few λ_W apart. This behavior is typical of a robust distance in statistics, meaning that outliers (over a few λ_W in distance) will have (almost) no effect on the optimization of the distance between the surfaces as it is an almost constant penalty.

In practice, surfaces are often represented by discrete triangulated meshes. Assuming that each face has a support which is smaller than ε times the scale λ_W , we can approximate the surface by the current $\mathcal{S} = \sum_k \delta_{x_k}^{\alpha_k}$, where x_k is the barycenter and α_k the normal weighted by the area of the face. From the above Taylor expansion of the distance, we can see that the approximation error is less than ε^2 for each face in the W^* norm.

The scalar product between two discrete currents is obtained by linearity:

$$\left\langle \sum_k \delta_{x_k}^{\alpha_k} \mid \sum_j \delta_{y_j}^{\beta_j} \right\rangle_{W^*} = \sum_{k,j} \alpha_k^\top K_W(x_k, y_j) \beta_j,$$

and the distance between two discrete surfaces is simply $\text{dist}^2(S, S') = \|\mathcal{S} - \mathcal{S}'\|_{W^*}^2$.

5.2.2.4 Effective Approximation of Currents using Matching Pursuit Algorithm

Representing shapes using currents gives a nice theoretical framework that allows us to compute simple statistics such as mean and principal modes. However, the complexity of the distance computation is quadratic in the number of Dirac delta currents used to represent the shape. This number can be quite high if we take highly detailed surface meshes, even though it adds nothing for the comparison of surfaces at a given scale λ_W . Therefore, in order to remain computationally as efficient as possible, we require a low number of degrees of freedom representation of currents that retains the information needed for the scale of the analysis. This is the essence of the matching pursuit algorithm as described in [24].

In brief, the matching pursuit algorithm is a greedy approach to find the best projection of multi-dimensional data on an over-complete family of basis functions [?]. In our case, we want to find an approximation of the current T that solves $\mathcal{L}_W^{-1}(T) = \gamma$, for a given vector field $\gamma \in W$. This amounts to finding N points (x_k) and vectors (α_k) such that the current $\Pi(T) = \sum_{k=1}^N \delta_{x_k}^{\alpha_k}$ is as close as possible to T . If the points are known, then $\Pi(T)$ is the orthogonal projection of T onto $\text{Span}(\delta_{x_k}^{\epsilon_q}; q = 1, 2, 3, k = 1 \dots N)$, where ϵ_q is the canonical basis of \mathbb{R}^3 . The orthogonality condition: $\langle T, \delta_{x_k}^{\epsilon_q} \rangle_{W^*} = \langle \Pi(T), \delta_{x_k}^{\epsilon_q} \rangle_{W^*}$ leads to a linear set of $3N$ equations

$$\sum_{p=1}^N (K_W(x_i, x_p) \alpha_p)_k = \gamma(x_i)_k, \quad (5.1)$$

which can be solved iteratively over the continuous space of Dirac delta currents to find the point positions (x_k) as well as their associated momenta and the residual vector field. Additionally, we can also sample the vector field on a grid Λ . Applying the matching pursuit in the discrete case constrains the estimated momenta to lie on the nodes of the grid which forces the estimated current to belong to a discrete set of currents W_Λ^* . This simplifies the problem in both complexity as well as computation time. Figure 5.4 shows a compressed representation of the right ventricle.




Fig. 5.4 Surfaces seen as integrators of vector fields. In a triangulated mesh, the atomic vector field integrators, the Dirac delta currents, are the normal vectors of every face, centered at the face barycenters and weighted by the face surface. Subfigure (a) displays the original mesh of the right ventricle with 1476 Diracs. In subfigure (b), a greedy algorithm reduces the amount of delta currents needed to represent the shape while preserving the accuracy of the representation to 281 Diracs, i.e., a compression of 80% with less than 5% error.

In this approximation algorithm, the interest is not to fix the number of points N and to look for the optimal representation (this cannot be achieved properly with a greedy algorithm), but rather to fix an error threshold (typically 1% to 5% of the norm of the original surface) and to stop the approximation whenever the number of points is sufficient to guaranty an error below that threshold. This typically results in compressions of 80% with less than a few percent of error for one surface, but compression by order of magnitudes for the mean of several currents.

5.2.3 An Algorithm for Surface Registration using Currents

Having defined a non-parametric representation of surfaces, we now require an algorithm to register one surface to another. Here, the term registration is meant to define the action of transforming two objects into the same co-ordinate system. The basic idea is to rephrase registration into a geodesic problem where the initial conditions are related to the formalism of currents.

Since we are comparing topologically similar shapes, the transformations are restricted to those which preserve the topology of the object and give a smooth one-to-one (invertible) transformation (i.e. a diffeomorphism). The space of diffeomorphisms give non-linear deformations that allow local smooth variations to be captured in the registration. However, diffeomorphic transformations have an infinite number of degrees of freedom, therefore optimizing over the whole group of diffeomorphisms may not be possible. So, we use a smaller infinite group of diffeomorphisms to allow computations with discrete parameterizations using the Large Deformation Diffeomorphic Metric Mapping (LDDMM) method described in [16, 17] for images.

The LDDMM framework uses a group of diffeomorphisms constructed through integration of time-varying vector fields that belong to a RKHS. This gives a geodesic flow of diffeomorphisms ϕ_t for a continuous parameter t within the interval $[0, 1]$. At time $t = 0$, we have the identity mapping ϕ_0 . The mapping at time 1 gives the desired transformation ϕ_1 which is required for mapping one image to

the other. The path of any point x is defined by $\phi_t(x)$ and leads to the final position $\phi(x) = \phi_1(x)$. By following the reverse path, we can compute the inverse deformation.

The basic idea of the LDDMM framework is to minimize the squared distance between objects after transformation (this is the similarity energy or data attachment term) with a penalization for the energy of the deformation trajectory (regularization). A time t , the speed at point $y = \phi_t(x)$ is $v_t(y) = d\phi_t(y)/dt$. This suggests to define the energy of a velocity field at deformation ϕ_t using a right-invariant metric:

$$\|v_t\|_{\phi_t}^2 = \|v_t \circ \phi_t^{(-1)}\|_V^2,$$

where $\|\cdot\|_V$ is the norm of the RKHS of the velocity fields at the identity transformation. The energy of the deformation trajectory is thus $E(\phi_t) = \int_0^1 \|v_t \circ \phi_t^{(-1)}\|_V^2 dt$, and the optimal curves joining $\phi_0 = Id$ to ϕ_1 are geodesics in the space of diffeomorphisms. Since optimal transformations are geodesics, we know that they are completely determined by their initial value, here the velocity field $v_0(x)$. As the velocity fields are most often non-zero everywhere, it is usually more efficient to parameterize the deformation with the the initial momentum $\mu_0 = K_V^{(-1)}v_0$, which is in most cases very sparse. Notice that the the momentum is an element of V^* , the dual space of V . Thus, it is a current, just as the ones we used to represent surfaces. However, even if the mathematical structures are identical, the objects “velocity fields” and “surfaces” are quite different and the kernels K_V and K_W are usually very different. Denoting ϕS the action of the transformation ϕ on the source object S , the registration criterion to the target object T is thus

$$C(\phi) = \text{dist}(\phi_1 S, T)^2 + \int_0^1 \|v_t \circ \phi_t^{(-1)}\|_V^2 dt.$$

For our surface registration problem, the objects are surfaces represented by their associated currents $\mathcal{S} = \sum_i \delta_{x_k}^{\alpha_k}$ and \mathcal{T} and the distance is taken in the space of currents W^* . In this setting, one can show that the optimal initial momentum of the deformation has the same point-wise support as \mathcal{S} : $\mu_0(x) = \sum_i \delta_{x_k}^{\beta_k}$ (i.e. $v_0(x) = \sum_i K_V(x, x_k)\beta_k$) for some set of vectors β_k . This means that we are left with a finite dimensional optimization problem [37, 7]. Notice that the number of unknowns is directly linked to the number of Diracs in the source surface \mathcal{S} , which is usually the atlas. This emphasizes the interest of compressing the atlas before registering it to target surfaces.

5.2.4 Building an Unbiased Atlas

When we have more than two surfaces, one has to choose a reference (called the atlas) to compare all the others consistently. Choosing this atlas randomly generally leads to biases in the forthcoming statistical analysis because the atlas shape might be far from the mean shape. The atlas is called unbiased if it is centered with respect to all the surfaces, i.e. if the mean deformation from the atlas to the observed surfaced is close to zero.

For building the optimally centered atlas, we took the ‘forward’ strategy that models the set of surfaces as the deformation of an unknown ‘ideal’ atlas plus some residuals [22, 23, 24]. For the patient number i , this can be expressed as

$$T_i = \phi_i \hat{T} + \varepsilon_i, \quad (5.2)$$

where \hat{T} is the atlas we are estimating, ϕ_i is the deformation that maps the atlas to the surfaces T_i , and ε_i represents the residuals (shape features not captured by the atlas such as changes in topology etc.). The mean shape information is described in the atlas \hat{T} while the shape variability is encoded in the transformation ϕ_i .

The atlas is first initialized by taking the mean of the patient meshes. This initial atlas is then registered to each of the patients individually. A new atlas \hat{T} that minimizes the error $\Lambda(\hat{T})$ is computed, with

$$\Lambda(\hat{T}) = \sum_{i=1}^M \|T_i - \phi_i(\hat{T})\|_{W^*}^2. \quad (5.3)$$

We then register the updated atlas to the individuals, recompute the atlas and loop until convergence (Algorithm 1).

Algorithm 1 Atlas Estimation

Input: M surface meshes T_i extracted from patient images.

- 1: Rigidly align meshes to a reference patient using rotations and translations.
 - 2: Create initial atlas \hat{T}_0 as the mean of the patient meshes (this is a current)
 - 3: **loop** {over k }
 - 4: Estimate the M transformations ϕ_i and their initial momenta μ_i that register the atlas \hat{T}_{k-1} to the individual T_i .
 - 5: Update the atlas by minimizing the error in Eq.5.3 using the estimated transformations ϕ_i and the atlas \hat{T}_{k-1}
 - 6: Compute bias $\bar{\mu} = \frac{1}{M} \sum_{i=1}^M \mu_i$ and variance $\sigma_\phi^2 = \frac{1}{M} \sum_{i=1}^M \|\mu_i\|_{V^*}^2$
 - 7: **end loop** [Stop whenever $\|\bar{\mu}\|_{V^*} < \eta \sigma_\phi$ where $\eta = 5\%$ typically]
 - 8: **return** Final atlas \hat{T}_k and the related transformations ϕ_i parameterized by their momenta μ_i
-

This algorithm produces an atlas which is a current and not any more a surface mesh. To constrain the atlas to remain a mesh, one usually restrains the search of the atlas to the subset of currents generated by all deformations of the data meshes. This means that the atlas can be written $\hat{T} = \phi T_l$ for some diffeomorphism ϕ and some index l . In practice, the atlas is initialized as the data mesh that has the smallest variance (sum of square distance to all other data meshes without deformation), and the update of the atlas becomes $\hat{T}_k = \bar{\mu} \hat{T}_{k-1}$.

5.3 Shape Analysis of ToF Data

To demonstrate the usefulness of our statistical shape analysis in a clinical context, we consider a population of patients with repaired tetralogy of Fallot. Recall that shape here means an atlas, which is a surface represented by a current, along with a deformation of that atlas. One of the key ideas of the ‘‘atlas framework’’ is to transfer the analysis from the space of currents (shapes) to the space of deformations.

As mentioned in the introduction, the clinical problem in rToF patients is in better understanding the shape re-modeling over time to aid in determining the optimal time for surgical intervention. In view of this problem, we are first interested in identifying the clinically relevant shape patterns by investigating the relationship between shape and given clinical indices. The assumption here is that changes in heart morphology may reveal structural abnormalities and dysfunctions due to the chronic regurgitation. In particular, we aim to establish which shape patterns are related to the pathology in order to give further insights into the condition. We can then go beyond the identification of pathological shape patterns, to estimate a predictive growth model of the heart to give an indication of how the heart will grow and re-model in time for this pathology.

Making use of currents to represent shapes, and an atlas to give an average representation of the population, we can apply classical statistical methods on the initial deformation momenta to address key questions for diagnosis, prognosis and prediction.

5.3.1 The Analysis Pipeline

Patient Data: We consider a data-set of thirteen patients (10 males, mean age \pm SD = 19 ± 9) with repaired tetralogy of Fallot. Steady-state Free Precision cine MRI (Fig. 5.5(a)) were acquired using a 1.5T scanner (Avanto, Siemens AG, Erlangen) in the short axis view covering entirely both ventricles (10-15 slices; isotropic in-plane resolution: 1.1×1.1 mm to 1.7×1.7 mm; slice thickness: 5mm to 10mm; 25-40 phases).

Surface Mesh Delineation: In the case of ToF patients, we are interested in studying the shape of the left and right ventricles of the heart. Using the statistical shape analysis methods described in




Fig. 5.5 Delineated boundaries of the right ventricle (pink), left ventricle endocardium (green) and left ventricle epicardium (yellow) shown on: (a) One image slice; (b) Two views of a 3D reconstruction.

the previous section, we can consider the ventricles as surfaces represented by triangulated meshes. The surface meshes of the right and left ventricle endocardium (inner layer of heart tissue), and left ventricle epicardium (outer layer of heart tissue) are defined using image segmentation by delineating the boundaries of each ventricle at end-diastole using the methods proposed by Zheng et al. [39] (Fig. 5.5). The right ventricle epicardium is not segmented because the thickness of the right ventricle is often not sufficient (with respect to the image resolution) to clearly distinguish the inner from the outer wall. This method defines the delineated boundaries using an anatomical model and thus establishes a point correspondence between meshes. Therefore, in this instance we used a standard least-squares method to first rigidly align the surface meshes to reduce the effect of patient positioning.

Mean Atlas Construction: Using the algorithm for computing a mean shape with the current method described in Sect. 5.2, we estimated an atlas for the ToF data, see Fig. 5.6. Eleven iterations of the alternate minimization for the shape atlas (Algorithm 1) were needed to reach convergence. The resulting atlas \hat{T} was well centered (mean over standard deviation of the deformations was 0.36946). Atlas-to-patient registration is implemented in parallel on a cluster of computers, which means that the atlas creation time is minimally dependent of the number of patients used to create it.

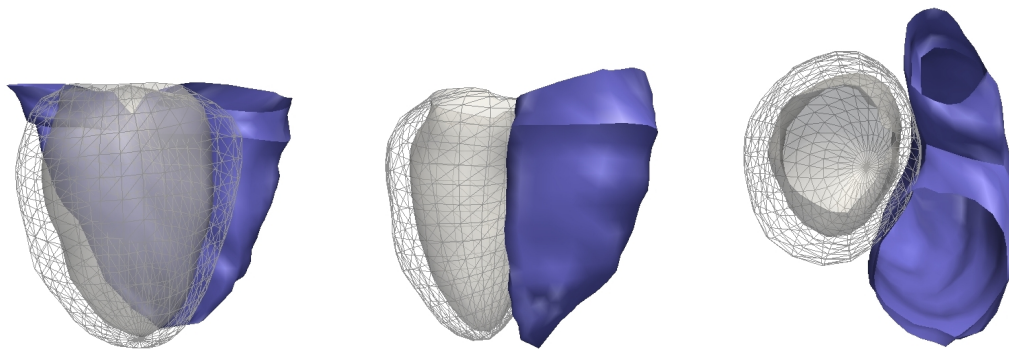


Fig. 5.6 Three views of the mean atlas estimated from 13 patients with repaired ToF with the right ventricle endocardium in blue, left ventricle endocardium in white and left ventricle epicardium in wire-frame.

5.3.2 Diagnosis Parameters

A key topic in computational shape analysis is identifying pathologically specific shape features in populations of diseased patients compared to controls, see [11] and references therein for instance. Beyond simply identifying the pathological shape features, we would like to also quantify the degree to which the shape is altered due to the pathology. One way in which this can be done is by correlating the shapes with clinical features to determine the severity of the disease, which requires a consistent representation of the patient shapes. For this we use firstly principal component analysis (PCA) to extract the main modes of shape variation followed by standard statistical design to exhibit those that are correlated to the pathology (Fig. 5.7 top row).

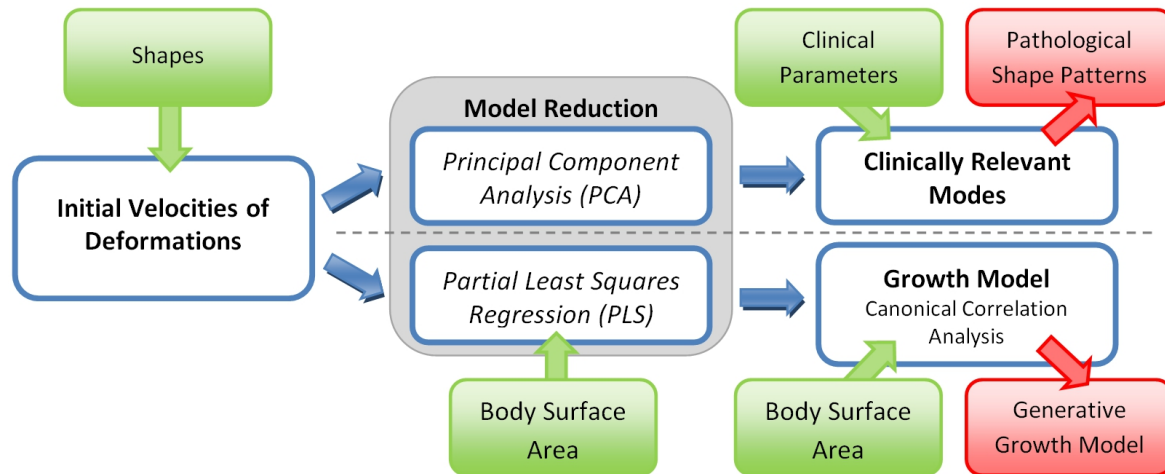


Fig. 5.7 Shape variation analysis pipeline with both clinically relevant model reduction and growth model generation. Using PCA modes and standard correlation analysis, the pathological shape patterns can be identified. Using a combination of PLS regression and CCA a statistical generative growth model can be derived.

5.3.2.1 Model Reduction using Principal Component Analysis

Since statistical shape analysis is a high dimensional problem with a large number of parameters and variables to solve for (despite the matching pursuit reduction, a shape can still be represented by hundreds of moments), we first reduce the dimension of the problem by applying principal component analysis (PCA). This gives the modes of deformation that describe the amount of variation of shape observed in the population.

PCA is applied on the initial velocity fields $v_0^{(i)}$ to extract the main deformation modes observed in the population. PCA finds basis vectors, the modes, of the space of variables (here the initial velocities) that best explain their variance. The modes p are calculated by solving the eigenvalue problem $\Sigma p = \mu p$, where the elements σ_{ij} of the covariance matrix Σ are calculated in the kernel space W . Assuming that the deformation from the atlas to patient i is parameterized by the initial vector field $v_0^{(i)}(x) = \sum_k K_W(x, x_k) \beta_k^{(i)}$, where the x_k are the point positions of the delta Dirac currents of the atlas, and $\beta_k^{(i)}$ the moment vector at x_k , then the mean initial vector field is $\bar{v}_0(x) = \sum_k K_W(x, x_k) \bar{\beta}_k$ and the covariance is

$$\sigma_{ij} = \langle v_0^i - \bar{v}_0, v_0^j - \bar{v}_0 \rangle_V = \sum_{k,l} (\beta_k^{(i)} - \bar{\beta}_k) K_W(x_k, x_l) (\beta_l^{(j)} - \bar{\beta}_l). \quad (5.4)$$

The principal components are obtained by computing the spectral decomposition $\Sigma = PMP^T$. M is the diagonal matrix of the eigenvalues σ_m , or variances, sorted in decreasing order and P is the orthonormal matrix (in the L_2 -norm sense) of the eigenvectors p^m . The m^{th} loading l^m of the PCA

decomposition is given by the formula:

$$l^m = \sum_i p^m[i] \beta^{(i)}. \quad (5.5)$$

In this equation, $p^m[i]$ is the i^{th} element of the m^{th} eigenvector of Σ , $\beta^{(i)} = (\beta_1^{(i)} - \bar{\beta}_1, \dots, \beta_n^{(i)} - \bar{\beta}_n)^T$ is the $n \times 3$ matrix that gathers the n centered moment of patient i . As a result, the initial velocity field of the m^{th} mode is $v_{0_{im}}(x) = \sum_k K_W(x, x_k) l^m$. The variability captured by this m^{th} deformation mode between $[-\eta\sigma^m; +\eta\sigma^m]$ is visualized by deforming the atlas T with the deformations ϕ^{-m} and ϕ^{+m} parameterized by the moments $\bar{\beta} - \eta l^m$ and $\bar{\beta} + \eta l^m$ respectively. Selecting the first p modes only among the $N - 1$ possible modes (where N is the number of patients) allows to explain a percentage $\sum_{m=1}^p \sigma^m / \text{trace}(\Sigma)$ of the total variance.

The orthogonal projection of each patient's initial velocity field onto the selected PCA subspace gives a unique shape vector. This simply corresponds to the coordinates of the projection in the basis constituted by the chosen eigenmodes:

$$s^{i,m} = \langle v_0^{(i)} - \bar{v}_0, v_{0_{im}} \rangle_W = \sum_{j,k} [\beta_j^{(i)} - \bar{\beta}] K_W(x_j, x_k) l_k^m. \quad (5.6)$$

Using PCA we have reduced the amount of data needed to represent the shape of a patient by two or more orders of magnitude. The precision of the representation is controlled by the number of components of the PCA subspace. However, it has to be observed that modes with low variances may still be relevant to external clinical parameters. For instance, a mode that captures a local bulging is probably more related to the pathology than a global scaling of the shape although this bulging is not very visible in the population and could be considered as noise in the model. Consequently, we are fairly conservative in the selection of the PCA subspace and select the modes based on their relationship with the clinical parameters of interest and not their variance, as described in the following sections.

5.3.2.2 Identifying Factors Between Shape and Clinical Features

The $s^{i,m}$'s quantify the amount of variability along the m^{th} mode present in patient deformation. We can thus investigate the heart shape by relating these shape vectors to clinical parameters that quantify the pathology. Ordinal clinical parameters are investigated using non-parametric rank-based statistics. Kruskal-Wallis analysis of variance is applied to find effects between the investigated parameters and shape [34]. If an effect is found, post-hoc two-sample Wilcoxon test is used to determine which levels differ [34]. Continuous clinical parameters are investigated using linear regression and Akaike Information Criterion (AIC) model reduction [1] to detect relevant modes and the direction of correlation.

To illustrate this method, in [33] we showed that these methods enable one to identify shape features related to the severity of the regurgitation for a data-set of 49 repaired ToF patients. The relationship between RV shape and pulmonary regurgitation were investigated by relating the PCA shape vectors with tricuspid regurgitation, trans-pulmonary valve regurgitation and pulmonary regurgitation volume indices taken from color Doppler ultrasound and phase-contract magnetic resonance images (PC-MRI). 90% of the spectral energy was explained by 18 PCA modes.

5.3.3 Building an Evolution Model

As explained beforehand, understanding and quantifying heart remodeling in these patients is crucial for planning pulmonary valve replacement. Given that there is a lack of longitudinal data available for these patients, we make use of the atlas as the mean of the population and cross-sectional statistics to formulate a generative growth model. Such a model could be used as reference, from which the pathology evolution of one patient could be quantitatively compared. In cross-sectional statistical

design, one does not propagate the evolution over time for a single patient but rather considers the image of each patient acquired at “time” t as an instance in the growth evolution of the group. In this way we can model the growth of the population given these instances using regression analysis (Fig. 5.8). By making use of the initial velocities v_0^i that parametrize the deformations ϕ^i computed using the methods described in the previous sections, we can regress the velocities against an index of patient growth using standard statistical techniques.

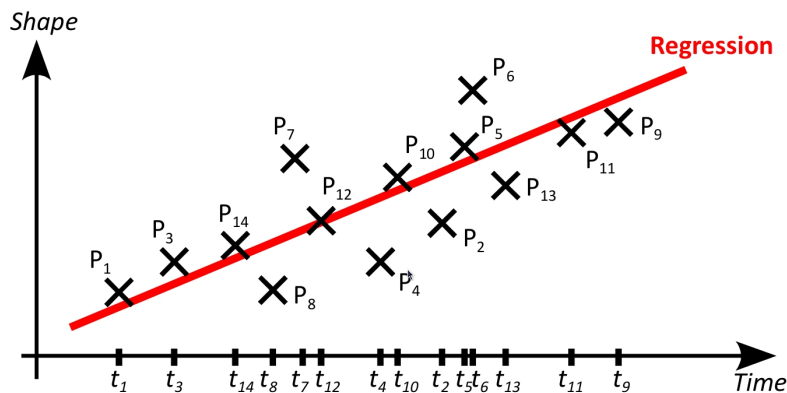


Fig. 5.8 Cross-sectional regression of shapes. Each patient is associated to a point in time (patient age for instance). A regression model is derived from the temporal data.

In order to obtain statistically significant results, we first need to reduce the dimensionality of the problem to consider just the factors related to patient growth, while also removing any co-linearity between factors. In the previous section the model reduction was performed using PCA. In this case we chose instead to use partial least squares (PLS) regression since it has the added advantage of computing the components that are most related to a given external parameter (i.e. patient growth). Using PLS allows us to compute the components that best describe the variance of both the matrix of predictors (X) and the matrix of responses (Y), as well as the covariance between X and Y , in a manner such that the regression $Y = f(X)$ is optimal.

In the case of ToF patients, we would ideally like to model the atlas deformation of a patient as a function of growth (i.e. $deformation = f(growth)$), however solving this problem is not possible due to the large number of deformation parameters that would need to be predicted with a single, one-dimensional, parameter (growth). Rather, we revert the problem to be a function of the deformations: $growth = f(deformations)$ which has a much lower number of parameters to predict. The output values are then projected onto the reduced PLS subspace and from that we can revert the problem to the desired form as a function of the shape using canonical correlation analysis (CCA) to give a generative growth model of the heart.

For this example we use body surface area (BSA) as the index of growth rather than patient age to better represent the growth given the variable age at which children enter puberty. We use the Dubois formula [5] to compute the BSA for each patient:

$$BSA(m^2) = 0.007184 \times weight(kg)^{0.425} \times height(cm)^{0.725}. \quad (5.7)$$

The pipeline we have just described for computing a generative growth model is shown in Fig. 5.7 (bottom row).

5.3.3.1 Model Reduction using Partial Least Squares Regression

As introduced in the previous section, regression using the partial least squares regression (PLS) method is based on finding an optimal basis of the predictor variables X that maximizes the variances of X and Y as well as their covariances. The method can be considered as the optimal estimation of two weight vectors r and s that satisfy

$$\max_{|r|=|s|=1} \text{cov}(Xr, Ys) = \max_{|r|=|s|=1} \text{var}(Xr) \text{corr}(Xr, Ys)^2 \text{var}(Ys), \quad (5.8)$$

under the constraint that the regression between X and Y is optimal. Mathematically, the generative model is

$$\begin{aligned} X &= \bar{X} + TP^T + E, \\ Y &= \bar{Y} + UQ^T + F, \end{aligned}$$

where T and U are the matrices of the PLS modes, P and Q are the loading matrices which describe the weight of each variable in $X - \bar{X}$ and $Y - \bar{Y}$ respectively and E and F are the residual terms which are the same size as X and Y respectively. Additionally the following regression condition is imposed for the PLS modes:

$$U = TD + G, \quad (5.9)$$

where D is a diagonal matrix of weights and G a matrix of residuals. Due to this added condition, the PLS loadings P and Q are not necessarily orthogonal as is the case for PCA modes.

Algorithm 2 Partial Least Squares Regression (PLS)

Input: Variables X and Y , number of components $p \leq N - 1$.

- 1: $X_1 = X - \bar{X}$, $Y_1 = Y - \bar{Y}$
 - 2: **for** $n = 1$ **to** p **do**
 - 2: $r_n \leftarrow$ first eigenvector of $X_n^T Y_n Y_n^T X_n$
 - 2: $t_n \leftarrow X_n r_n / \|r_n\|$ n^{th} PLS component of X
 - 2: $s_n \leftarrow Y_n t_n / (t_n^T t_n)$
 - 2: $u_n \leftarrow Y_n s_n / \|s_n\|$ n^{th} PLS component of Y
 - 2: $p_n \leftarrow X_n t_n / (t_n^T t_n)$ n^{th} loading of X
 - 2: $q_n \leftarrow Y_n u_n / (u_n^T u_n)$ n^{th} loading of Y
 - 2: $X_{n+1} \leftarrow X_n - t_n p_n^T$ deflation of X
 - 2: $Y_{n+1} \leftarrow Y_n - t_n [t_n^T Y_n / t_n^T t_n]$ deflation of Y
 - 3: **end for**
 - 4: **return** $T = (t_n)_{n=1\dots p}$, $P = (p_n)_{n=1\dots p}$, $U = (u_n)_{n=1\dots p}$, $Q = (q_n)_{n=1\dots p}$
-

Several algorithms have been proposed to compute the PLS modes. In this work, we use the PLS1 method, an efficient iterative algorithm that does not require matrix inversion as summarized in Algorithm 2. X is the matrix of the initial velocity field moments for all patients and Y is the vector of the BSA values for all patients. The first five PLS modes are shown in Fig. 5.9 and the explained variance and correlation of modes is shown in Fig. 5.10. These modes account for 99% of the covariance between shape and BSA in the population and 61% of the shape variability and are oriented along increasing BSA. Visually we can see that Modes 1 and 3 display an overall dilation in both the left and right ventricles. The second mode shows a narrowing in the right ventricular outflow tract with a noticeable dilation in the left ventricle which can be seen in the top view of the mode. Modes 4 and 5 show an elongation at the right ventricular outflow tract. Mode 5 also shows an aneurysm in the right ventricle outflow tract which is characteristic of rToF patients, although it is still not clear if this is due to pulmonary regurgitations (although both are correlated) or the initial surgical patch.

5.3.3.2 Generating a Growth Model using Canonical Correlation Analysis

Using PLS as described above allows us to predict BSA given the shape, however what we would like is to estimate the shape given BSA. To reverse the relationship we use canonical correlation analysis (CCA) on the PLS shape vectors (the t_i 's of Algorithm 2). CCA computes the vectors r and s that maximize the correlation between the two sets X and Y

$$\max_{|r|=|s|=1} \text{corr}(Xr, Ys)^2, \quad (5.10)$$

where Y is the vector of BSA values and $X = [t_1, \dots, t_N]^T$ is the matrix of shape descriptors.

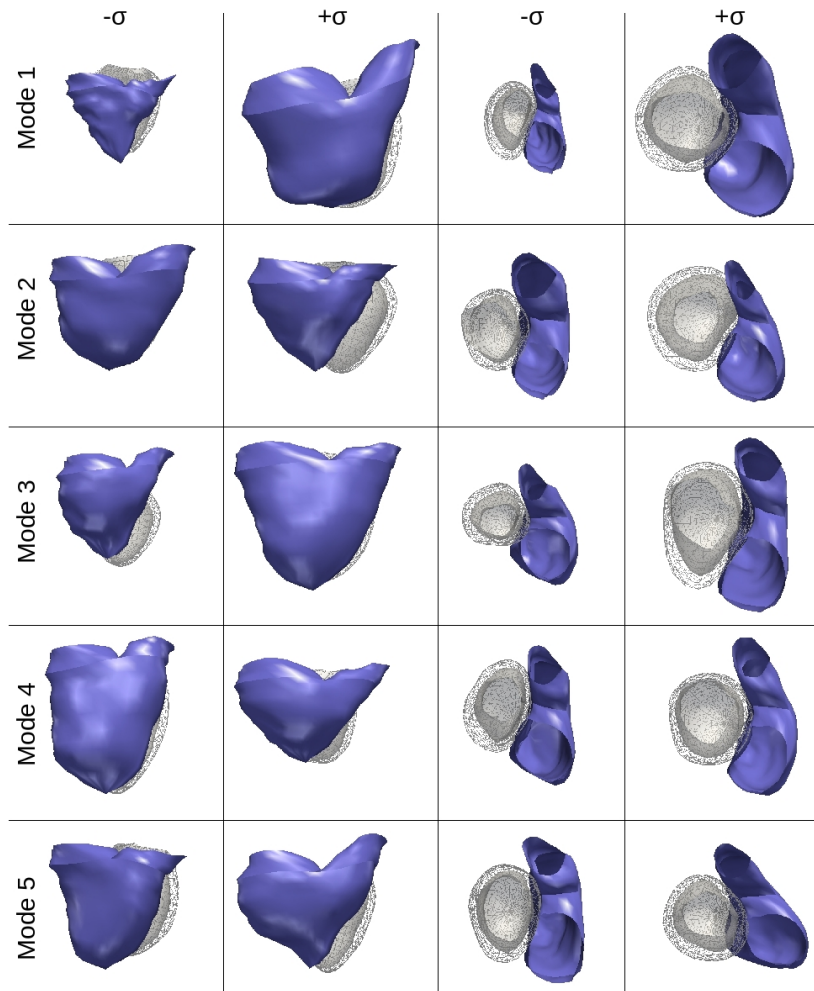


Fig. 5.9 The first five PLS modes of variation that describe 99% of the observed BSA variability and 61% of the observed shape variability in the population.

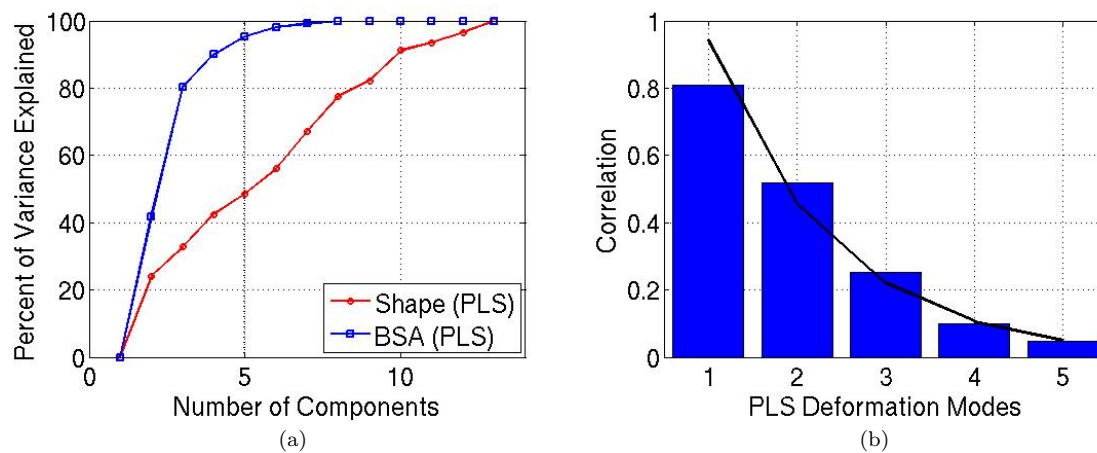


Fig. 5.10 Analysis of the variance and dependency of the shape and covariates using PLS. (a) Cumulative variance of PLS modes with respect to shape and BSA. (b) CCA correlation coefficients between BSA and PLS modes with exponential fitted curve in black.

If we define the covariance matrices $V_{UZ} = \frac{1}{N-1}U^T Z$, then the matrix $\Gamma = V_{XX}^{-1/2}V_{XY}V_{YY}^{-1/2}$ can be seen as a multi-variate generalization of the uni-dimensional correlation coefficient $r = \sigma_{XY}/\sqrt{\sigma_{XX}\sigma_{YY}}$. The sought correlations are obtained by SVD decomposition of Γ :

$$\Gamma = ASB^T, \quad (5.11)$$

where S is the diagonal matrix of the correlation coefficients between correlation vectors and A and B are rotation matrices of correlation vectors, i.e., $A^T A = B^T B = Id$. In our application, Y is a one-column matrix. Hence, S has only one non-null coefficient R , which is the overall correlation between the PLS shape vectors X and BSA. B is a scalar equal to ± 1 that determines the direction of BSA correlation. The elements of the first correlation vector of A , denoted by ρ , relate to the amplitude and direction of correlations of each predictor, namely each PLS mode, when Y varies along the direction defined by the sign of B . In other words, when BSA varies by 1, the k^{th} predictor varies by $BR\rho[k]$. We can therefore compute a generative average model of heart growth by artificially increasing BSA and deforming the atlas T with the growth deformation Φ parametrized by the moments $\mu = B.R.\sum_k \rho[k]p^k$, where p^k is the k^{th} PLS loading.

5.3.3.3 Interpretation

The growth model computed on the rToF data-set is shown in Fig. 5.11. This model shows an expected overall growth of both ventricles as body surface area increases. We can also see the caving of the septum into the right ventricle as time passes and the elongation of the right ventricular outflow tract which is observed in these patients over time.

Given a larger data-set this model can be improved further by increasing the certainty in the model and by capturing more variability in shape observed in a wider population. As well, given more data we can divide the patients according to an external parameter such as treatment group to formulate a growth model for a given course of treatment. In the case of ToF, this could allow clinicians to have a better idea of how the heart re-models after different types of valve replacement surgery and more importantly, the effect of the initial surgery on the long-term outcome. This is the key question.




Fig. 5.11 Mean growth model computed from a population of 13 rToF patients. Both ventricles grow as BSA increases

5.4 Conclusion

In this chapter we presented methods for computing statistics on shapes. The proposed methods rely on currents to represent the anatomical shapes in a consistent way without the need for defining landmarks or point correspondences between shapes. This allows the computation of population statistics such as the mean and variance observed within a population and to correlate shape with clinical indices, for instance to quantify the severity of a disease. We also present a framework for formulating a generative statistical growth model to simulate the growth of an organ over time. This is especially important for our clinical application where the shape remodeling of the heart in congenital diseases like Tetralogy of Fallot is crucial to better understand the evolution of the disease and ultimately to aid for therapy planning.

From a clinical perspective, most of the shape changes that could be put into evidence in our longitudinal study were corresponding to some previously reported observations in the literature. However, we have for the first time a combined and quantitative model of all the changes together. Of course, the number of parameters that could influence the shape of the heart is so large that we would need much more subjects in order to obtain an accurate and reliable quantitative model of the heart remodeling. Nevertheless, the shape analysis pipeline we proposed is more than a black box that learns the statistical relationship between sets of variable: it gives a generative shape model that can be explored. For instance, the visualization of the remodeling (or the deformation correlated to other clinical variables) could allow to draw some intuition about the potential underlying links. Ideally this would lead to a simplified model that could be tested against new data, thus gradually improving the knowledge that we have about the disease. As the anatomy is the support of the physiology, one of the goals is now to couple the anatomical changes with the physiological evolution in order to better understand how one influences the other.

Other interesting clinical questions related to rToF could be investigated using our approach. For instance, there is nowadays a clinical consensus that a longer QRS complex duration measured on the electrocardiogram (see for instance Fig. ?? in chapter ??) correlates with RV dilation. Correlating this parameter with the RV shape could reveal how the abnormal conduction impacts the RV anatomy and function. Similarly, studying body-mass index (BMI) jointly with BSA could provide a more comprehensive representation of the patient growth. One could apply the PLS method between two sets of multivariate variables (BMI, BSA and QRS on the one hand, the shape vectors on the other hand) and get a growth model that explains both features. Unfortunately, QRS durations and BMI were not available for all the patients as the population was retrospective, keeping these questions for future work. The effect of genetic factors that regulate myocardium stiffness on the long-term RV remodeling could also be analyzed. Patients with stiffer myocardium are known to be more protected against RV dilation. These patients may be the outliers of the model with “abnormally” smaller RV than the average. The decision for valve implant may be delayed and based on different features for these patients. Finally, tools are now been developed and integrated in several hospitals to keep track of the patient records during a long time-period, which opens the possibility to now correlate the shape features directly with the clinical outcome. Such a shape-based model of the outcome, when validated, would be immediately usable by the clinicians as a computer aided diagnosis tool.

5.5 Online Resources

[Integrated Models of Paediatric Heart Diseases](http://www-sop.inria.fr/asclepios/projects/Health-e-Child/DiseaseModels/content/cardiac/index.html)

<http://www-sop.inria.fr/asclepios/projects/Health-e-Child/DiseaseModels/content/cardiac/index.html>

This site summarizes some modeling issues in pediatric cardiac diseases. It explains the approach that was adopted in the European FP6 Health-e-Child project.

[Shape analysis software: Exoshape](http://www-sop.inria.fr/asclepios/projects/Health-e-Child/ShapeAnalysis/)

<http://www-sop.inria.fr/asclepios/projects/Health-e-Child/ShapeAnalysis/>

This web site contains the source of the Matlab software developed for doing the analysis described in this chapter, and a description of its application to the analysis of the shape right ventricle of the heart in ToF.

The Virtual Physiological Human

<http://ercim-news.ercim.eu/en69/special>

April 2007 issue of the ERCIM News journal on “The Digital Patient” describing many different aspects of the Virtual Physiological Human initiative to develop in silico models of human physiology and pathology.

http://en.wikipedia.org/wiki/Virtual_Physiological_Human

The wikipedia description of the VPH with links to other resources

Acknowledgment

The computational tools used in this chapter were originally developed within the context of the European FP6 project Health-e-Child (<http://www.health-e-child.org/>). The software was made available to the community in collaboration with the EU network of Excellence Virtual Physiological Human (<http://www.vph-noe.eu/>). The extension to the analysis of the bi-ventricular shape of the heart in rToF was performed in the context of the European ITEA2 Care4Me project (www.care4me.eu/).

List of Acronyms

BMI	Body Mass Index
BSA	Body Surface Area
CCA	Canonical Correlation Analysis
LDDMM	Large Deformations Diffeomorphic Metric Mapping
MRI	Magnetic Resonance Imaging
PCA	Principal Component Analysis
PLS	Partial Least Squares
RKHS	Reproducing Kernel Hilbert Space
rToF	repaired Tetralogy of Fallot
RV	Right Ventricle
ToF	Tetralogy of Fallot

References

1. H. Akaike. A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6):716723, 1974.
2. F.L. Bookstein. *The Measurement of Biological Shape and Shape Change*, volume 24 of *Lecture Notes in Biomathematics*. Springer-Verlag, 1978.
3. F.L. Bookstein. Size and shape spaces for landmark data in two dimensions (with discussion). *Statistical Science*, 1:181–242, 1986.
4. I.L. Dryden and K.V. Mardia. Theoretical and distributional aspects of shape analysis. In *Probability Measures on Groups, X (Oberwolfach, 1990)*, pages 95–116, New York, 1991. Plenum.
5. D. DuBois and E. DuBois. The measurement of the surface area of man. *Archives of Internal Medicine*, 15(5), 1915.
6. S. Durrleman. *Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution*. Thèse de sciences (phd thesis), Université de Nice-Sophia Antipolis, 2010.
7. J. Glaunès. *Transport par difféomorphismes de points, de mesures et de courants pour la comparaison de formes et anatomie numérique*. Thèse de sciences, Université Paris 13, November 2005.
8. U. Grenander. *General Pattern Theory: A Mathematical Study of Regular Structures*. Oxford University Press Inc., New York, NY., 1993.
9. J. Hoffman and S. Kaplan. The incidence of congenital heart disease. *Journal of the American College of Cardiology*, 39(12):1890–1900, 2002.
10. H. Hufnagel. *A probabilistic framework for point-based shape modeling in medical image analysis*. Phd thesis, University of Lübeck, 2010.

11. P. Fletcher J. Cates and R. Whitaker. A hypothesis testing framework for high-dimensional shape models. In *In MICCAI Workshop on Mathematical Foundations of Computational Anatomy*, page 170, 2008.
12. D.G. Kendall. A survey of the statistical theory of shape (with discussion). *Statistical Science*, 4:87–120, 1989.
13. H. Le and D.G. Kendall. The Riemannian structure of Euclidean shape space: a novel environment for statistics. *Annals of Statistics*, 21:1225–1271, 1993.
14. L. Younes M. Vaillant, M.I. Miller and A. Trouvé. Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, 23(Supp. 1):S161–S169, 2004.
15. T. Mansi. *Image-Based Physiological and Statistical Models of the Heart, Application to Tetralogy of Fallot*. Thèse de sciences (phd thesis), Ecole Nationale Supérieure des Mines de Paris, 2010.
16. A. Trouvé M.F. Beg, M.I. Miller and L. Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. Journal of Computer Vision*, 61(2):139–157, 2005.
17. A. Trouvé M.I. Miller and L. Younes. On the metrics and Euler-Lagrange equations of computational anatomy. *Annual Re-view of Biomedical Engineering*, pages 375–405, 2003.
18. M.I. Miller and L. Younes. Group actions, homeomorphisms, and matching: A general framework. *International Journal of Computer Vision*, 41(1/2):61–84, 2001.
19. J. T. Ratnanather R. A. Poldrack Th. E. Nichols J. E. Taylor P. M. Thompson, M. I. Miller and K. J. Worsley, editors. *Mathematics in Brain Imaging*, volume 45 of *NeuroImage, special issue*. Springer, March 2004.
20. X. Pennec. Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, July 2006. A preliminary appeared as INRIA RR-5093, January 2004.
21. X. Pennec. Statistical computing on manifolds: from riemannian geometry to computational anatomy. In *Emerging Trends in Visual Computing*, volume 5416 of *LNCS*, pages 347–386. Springer, 2008.
22. Y. Amit S. Allassonnire and A. Trouv. Towards a coherent statistical framework for dense deformable template estimation. *Journal Of The Royal Statistical Society Series B*, 69(1):3–29, 2007.
23. A. Trouvé S. Durrleman, X. Pennec and N. Ayache. A forward model to build unbiased atlases from curves and surfaces. In X. Pennec and S. Joshi, editors, *Proc. of the International Workshop on the Mathematical Foundations of Computational Anatomy (MFCA-2008)*, September 2008.
24. A. Trouvé S. Durrleman, X. Pennec and N. Ayache. Statistical models on sets of curves and surfaces based on currents. *Medical Image Analysis*, 13(5):793–808, 2009.
25. A. Trouvé G. Gerig S. Durrleman, X. Pennec and N. Ayache. Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI'09), Part I*, volume 5761 of *Lecture Notes in Computer Science*, pages 297–304, London, UK, 2009. Springer.
26. A. Trouvé P. Thompson S. Durrleman, X. Pennec and N. Ayache. Inferring brain variability from diffeomorphic deformations of currents: an integrative approach. *Medical Image Analysis*, 12(5):626–637, 2008.
27. A. Srivastava S. H. Joshi, D. Kaziska and W. Mio. Riemannian structures on shape spaces: A framework for statistical inferences. In Hamid Krim and Anthony Yezzi, editors, *Statistics and Analysis of Shapes*, Modeling and Simulation in Science, Engineering and Technology, pages 313–333. Birkhuser Boston, 2006.
28. A. Srivastava S. H. Joshi, E. Klassen and I. Jermyn. A novel representation for riemannian analysis of elastic curves in rn. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 0:1–7, 2007.
29. A. Srivastava S. H. Joshi, E. Klassen and I. Jermyn. Removing shape-preserving transformations in square-root elastic (sre) framework for shape analysis of curves. In *EMMCVPR'07*, pages 387–398, 2007.
30. M. Jomier S. Joshi, B. Davis and G. Gerig. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage*, 23(Supplement 1):S151 – S160, 2004. Mathematics in Brain Imaging.
31. C.G. Small. *The Statistical Theory of Shapes*. Springer series in statistics. Springer, 1996.
32. B. Bernhardt M. Sermesant H. Delingette I. Voigt T. Mansi, S. Durrleman, J. Blanc Y. Boudjemline X. Pennec P. Lurz, A. M. Taylor, and N. Ayache. A statistical model of right ventricle in tetralogy of fallot for prediction of remodelling and therapy planning. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'09)*, volume 5761 of *Lecture Notes in Computer Science*, pages 214–221. Springer, 2009.
33. B. Leonardi X. Pennec S. Durrleman M. Sermesant H. Delingette A. M. Taylor Y. Boudjemline G. Pongiglione T. Mansi, I. Voigt and N. Ayache. A statistical model for quantification and prediction of cardiac remodelling: Application to tetralogy of fallot. *IEEE Transactions on Medical Images*, 2011.
34. R Development Core Team. R: A language and environment for statistical computing. *R Foundation for Statistical Computing, Vienna, Austria*, Online. Available: <http://www.R-project.org>, 2009.
35. W. D'Arcy Thompson. *On Growth and Form*. Cambridge University Press, England., 1917.
36. A. Trouvé. Diffeomorphisms groups and pattern matching in image analysis. *International Journal of Computer Vision*, 28(3):213–221, 1998.
37. M. Vaillant and J. Glaunes. Surface matching via currents. In *Proc. of IPMI'05*, pages 381–392, 2005.
38. A. Srivastava W. Mio and S. H. Joshi. On shape of plane elastic curves. *International Journal of Computer Vision*, pages 307–324, 2007.
39. B. Georgescu M. Scheuering Y. Zheng, A. Barbu and D. Comaniciu. Four-chamber heart modeling and automatic segmentation for 3-d cardiac ct volumes using marginal space learning and steerable features. *IEEE Trans. on Medical Imaging*, 27(11):16681681, 2008.