



HAL
open science

Parallel transport, a central tool in geometric statistics for computational anatomy: Application to cardiac motion modeling

Nicolas Guigui, Xavier Pennec

► To cite this version:

Nicolas Guigui, Xavier Pennec. Parallel transport, a central tool in geometric statistics for computational anatomy: Application to cardiac motion modeling. Frank Nielsen; Arni S.R. Srinivasa Rao; C.R. Rao. *Geometry and Statistics*, 46, Elsevier, pp.285-326, 2022, *Handbook of Statistics*, 978-0-323-91345-4. 10.1016/bs.host.2022.03.006 . hal-03684811

HAL Id: hal-03684811

<https://inria.hal.science/hal-03684811v1>

Submitted on 1 Jun 2022

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.

Parallel transport, a central tool in geometric statistics for computational anatomy. Application to cardiac motion modelling

Nicolas Guigui¹ Xavier Pennec¹

¹Université Côte d’Azur and INRIA, Epione team, Sophia Antipolis, France

February 2022

Contents

1	Introduction	2
1.1	Diffeomorphometry	3
1.2	Longitudinal models	4
1.3	Parallel transport for inter-subject normalisation	6
1.4	Chapter organization	8
2	Parallel transport with ladder methods	9
2.1	Numerical accuracy of Schild’s and pole ladders	10
2.2	A short overview of the LDDMM framework	13
2.3	Ladder methods with LDDMM	15
3	Application to cardiac motion modelling	18
3.1	The Right Ventricle and its diseases	18
3.2	Motion normalization with parallel transport	19
3.3	An intuitive rescaling of LDDMM parallel transport	22
3.4	Changing the metric to preserve relative volume changes	25
3.5	Analysis of the normalized deformations	28
3.5.1	Geodesic and Spline regression	28
3.5.2	Hotelling tests on velocities	29
4	Conclusion	31

Abstract

Transporting the statistical knowledge regressed in the neighbourhood of a point to a different but related place (transfer learning) is important for many applications. In medical imaging, cardiac motion modelling and structural brain changes are two such examples: for a group-wise statistical analysis, subject-specific longitudinal deformations need to be transported in a common template anatomy.

In geometric statistics, the natural (parallel) transport method is defined by the integration of a Riemannian connection which specifies how tangent vectors

are compared at neighbouring points. In this process, the numerical accuracy of the transport method is critical. Discrete methods based on iterated geodesic parallelograms inspired by Schild’s ladder were shown to be very efficient and apparently stable in practice. In this chapter, we show that ladder methods are actually second order schemes, even with numerically approximated geodesics. We also propose a new original algorithm to implement these methods in the context of the Large Deformation Diffeomorphic Metric Mapping (LDDMM) framework that endows the space of diffeomorphisms with a right-invariant RKHS metric.

When applied to the motion modelling of the cardiac right ventricle under pressure or volume overload, the method however exhibits unexpected effects in the presence of very large volume differences between subjects. We first investigate an intuitive rescaling of the modulus after parallel transport to preserve the ejection fraction. The surprisingly simple scaling/volume relationship that we obtain suggests to decouple the volume change from the deformation directly within the LDDMM metric. The parallel transport of cardiac trajectories with this new metric now reveals statistical insights into the dynamics of each disease. This example shows that parallel transport could become a tool of choice for data-driven metric optimization.

Keywords: Parallel transport, longitudinal studies, mean trajectory, cardiac motion analysis, Schild’s ladder, pole ladder, Riemannian manifolds.

1 Introduction

At the interface of geometry, statistics, image analysis and medicine, *Computational Anatomy* aims at analysing and modelling the variability of the biological shape of tissues and organs and their dynamics at the population level. The goal is to estimate representative anatomies across diseases, populations, species or ages, to model the organ development across time (growth or ageing), to discover morphological differences between normal and pathological groups and to estimate and correlate the variability with other functional, genetic or structural information. In the context of cardiology, computational anatomy models the cardiac cycle as a sequence of deformations of an anatomy and allows to characterize quantitatively these deformations to compare the impact of diseases on the cardiac function.

The analysis of the organs’ shape and deformations often relies on the identification of features to describe locally the anatomy such as landmarks, curves, surfaces, intensity patches, full images, etc. Modelling their statistical distribution in the population requires to first identify point-to-point anatomical correspondences between these geometric features across subjects. This may be feasible for landmark points, but not for curves or surfaces. Thus, one generally considers relabelled point-sets or reparametrized curves/surfaces/images as equivalent objects. With this geometric formulation, shape spaces are the quotient of the original space of features by their reparametrization group. The difficulty is that taking the quotient generally endows the shape space with a non-linear manifold structure, even when we start from features living in a Euclidean space.

For instance, equivalence classes of k -tuples of points under rigid or similarity transformations result in non-linear Kendall’s shape spaces (see e.g. Dryden and Mardia 2016 for a recent account on that subject). The quotient of curves, surfaces

and higher dimensional objects by their reparametrizations (diffeomorphisms of their domains) produces in general even more complex infinite dimensional shape spaces (Bauer, Bruveris, and Michor 2014). Thus, shapes belong in general to non-linear manifolds, while statistics were essentially developed for linear and Euclidean spaces. This has motivated the use and development of a consistent statistical framework on Riemannian manifolds and Lie groups during the past 25 years, a field called *Geometric Statistics* in computational anatomy (Pennec, Sommer, and Fletcher 2020).

Deformation-based morphometry (DBM) is also a popular method to study statistical differences of brain anatomies. It consists in analyzing local deformation features of non-linear image registration to a reference (Ashburner et al. 1998). For instance, the Jacobian of the transformation encoding local volumes changes can be used to detect the areas that are statistically growing or shrinking in a population. However, DBM analyses the deformation independently at each point of the image. In order to capture spatial correlations, it is more interesting to model the transformations of all the points together, i.e. to lift statistics from the objects (image voxels, curves, surfaces) to the transformations of their embedding space. This powerful idea also allows us to capture jointly the variability of several structures to model their interactions. In the brain, for instance, one can consider together the shape of the cortex, the ventricles, deep grey nuclei and other internal brain structures. We can also include structures that are not surface, such as the sulcal lines encoding the complex folding patterns of the grey matter at the surface of the cortex.

1.1 Diffeomorphometry

Following D’Arcy Thompson 1917, it is often assumed that there exists a template shape or image (called an atlas in medical image analysis) that represents the standard (or mean) anatomy, and that the inter-subject variability is encoded by deformations of that template towards the shapes of each subject or their evolution in time. A very desirable aspect of these transformations is to smoothly preserve the spatial organization of the anatomical tissues by avoiding intersections, folds or tears. Simply encoding deformations with the vector space of displacement fields is not sufficient to preserve the topology: one needs to require diffeomorphic transformations (differentiable one-to-one transformations with differentiable inverse).

Lie groups of diffeomorphisms are examples of space that are both infinite dimensional manifolds and Lie groups, and the statistic analyse of shapes through their diffeomorphic deformations has been coined as *diffeomorphometry*. This approach was pushed forward by (Grenander and M. Miller 1998) and turned into a mathematically grounded framework by endowing the space of diffeomorphisms with a sufficiently regular right invariant metric (Younes 2019; M. Miller, Trounev, and Younes 2015), leading to the so called *Large Deformation Diffeomorphic Metric Mapping (LDDMM)* framework, detailed in Section 2.2.

Since the optimization of time-varying velocity fields was computationally intensive, an alternative parameterization of diffeomorphisms with the flow of stationary velocity fields (SVF) was introduced by Arsigny et al. 2006. The flow of SVFs generates one-parameter subgroups, which are simply matrix exponentials in matrix Lie groups. Although a number of theoretical difficulties remain when moving to infinite dimensions, very efficient algorithms can be adapted from the matrix case, like the

scaling and squaring procedure (Higham 2005) to compute the group exponential and its Jacobian, or the Baker-Campbell-Hausdorff (BCH) formula to approximate the composition of two deformations directly in the log-domain. This allows the straightforward extension of the classical “demons” image registration algorithm to encode diffeomorphic transformations parametrized by SVFs. Because the inverse of the flow of an SVF is simply parameterized by the opposite of this SVF, a special feature of the log-demons registration framework is to enforce almost seamlessly the very desirable inverse consistency of the registration. The SVF framework was successfully used in a number of application to brain studies (Lorenzi et al. 2011; Hadj-Hamou et al. 2016) as we will see below.

The differential geometric foundations of this very efficient SVF framework were uncovered in (Lorenzi and Pennec 2013): one-parameter subgroups are actually geodesics (in the sense of auto-parallel curves) of the Cartan-Schouten connection, a canonical bi-invariant symmetric affine connection that exists on every Lie group. Whenever a bi-invariant metric exists on the group, this affine connection coincides with the Levi-Civita connection of that bi-invariant metric, in which case the Fréchet mean is left, right and inverse equivariant. However, only direct products of compact and Abelian groups admit a bi-invariant metric, which is very limiting in practice. Thus, Lie groups generally do not admit any bi-invariant metric, in which case the Fréchet mean based on a left-invariant distance is not consistent with right translations nor with inversions. This is already the case for Euclidean motions $SE(3)$ since it is a semi-direct product. In contrast, exponential barycenters of the canonical Cartan-Schouten connection define a non-metric notion of mean on any Lie groups which is automatically equivariant by left, right translation and inversion (Pennec and Arsigny 2012). One can also define covariance matrices and Mahalanobis distances that are consistent with the group operations. Thus, from the points of view of abstract transformations, independently of their action on objects, the Cartan-Schouten connection offers a unique way to define a consistent statistical framework on Lie groups. This drove the interest for statistics on affine connection spaces, a superclass of Riemannian manifolds where non-metric geodesics are defined locally as auto-parallel curves (Pennec and Lorenzi 2020).

1.2 Longitudinal models

Once we have decided for the LDDMM or the SVF parameterization of diffeomorphisms, a first and simple approach to tackle longitudinal data is to regress the template and its deformations over time by minimizing the inter-subject registration square-distance to each observation. The template trajectory may be a geodesic in a group of diffeomorphism for very few data points, such as for studying the remodelling of the heart over time in paediatric images (Mansi et al. 2009; McLeod et al. 2013). With more samples over time, one may fit a deformation trajectory with kernel regression (Gerber et al. 2010) or assume a spline structure (Trouvé and Vialard 2012b; Hinkle, Fletcher, and Joshi 2014; Singh, Vialard, and Niethammer 2015). This type of method is good for a cross-sectional design, where we have one data point per subject, but is likely to hide the small longitudinal changes of each subject into the large inter-subject variability when we have more than one time-point per subject.

Longitudinal models are preferable to study time-series of shapes in both dynam-

ical systems such as the heart and disease progression modelling such as the aging brain. Due to the non-commutativity of the longitudinal and inter-subject deformations, several paradigms were investigated to compare transformations trajectories at different points (Durrleman et al. 2013). One choice is to model the reference trajectory in the template space and to make it patient-specific by deforming it using a template to subject transformation valid for all times (Rao et al. 2004; Durrleman et al. 2013). This approach can be geometrized by parallel transporting the initial velocity of the template to subject deformation along the reference longitudinal deformation, as proposed by (Schiratti et al. 2015). This was called exp-parallelization for a geodesic reference trajectory, but the method works for any smooth curve. A recent refinement of this model has been proposed to cluster trajectories into distinct classes with different trajectories (Debavelaere et al. 2019). However, this does not allow to easily model the variability on the longitudinal trajectories which is observed for instance in the heart. Moreover, exp-parallel curve may be very different from geodesics when the atlas-to-subject transformation is large: on a sphere, for instance, exp-parallel curves are small circles parallel to the reference geodesic great circle, and reduce to a point at the poles.

In this chapter, we prefer to use the other main approach, where we first regress a patient-specific trajectory, and then normalize it with respect to the patient-to-atlas deformation at a reference time to perform the statistical analysis of all trajectories at the population level (Lorenzi and Pennec 2013). It is sometimes argued that there is no canonical reference time for comparing the evolution of brains. This is not the case for the heart, for which the end-diastole (ED) is a natural time reference of the cardiac cycle. The subject-specific trajectory is usually a geodesic for the analysis of structural brain changes with ageing in Alzheimer’s disease (Hadj-Hamou et al. 2016). For the heart, one can use any motion trajectory parametrized by its logarithm (the initial tangent vector or momentum that geodesically shoots to the deformation at a given point of the trajectory). In this setting, one has a curve of tangent vectors encoding the subject-specific deformation over time that needs to be transported to a reference template.

Transporting scalar values (e.g. the intensity of an image) along a diffeomorphic transformation is simply done by resampling that image. However, when one wants to transport differential (vector) quantities, it is not sufficient to resample each of the vector component individually: there needs to be a normalization that also reorient the vector. For instance, a constant initial vector field encodes a translation over time. However, when the inter-subject transformation is a rotation by 90 degrees, that vector needs to be rotated along. In this chapter, we focus on geometric methods based on parallel transport.

For cardiac motion, an average cardiac deformations across patients is called a *statistical motion atlas* (Young and Frangi 2009; Peyrat et al. 2010; Duchateau et al. 2011). In this application, the need to normalize deformations is salient in the example of Figure 1, where the deformation between the patients’ end-diastolic (ED) and end-systolic (ES) frames are applied without any normalization to another reference shape, the atlas. These result in meshes of poor quality, with irregularities, especially near the valves and on the septum. These abnormalities are salient when considering ventricles with pressure and volume overload as the amplitude of both the intra-patient and patient-to-atlas deformations may be large.




Figure 1: Example of a systolic deformation estimated on a patient and applied to the atlas without normalization (middle) and with normalization by the method of Section 3.3 (right). The ED frame is a red point cloud, while the ES frame is the blue mesh. The ES frames obtained without normalization show irregularities and cannot be used in downstream analyses.

1.3 Parallel transport for inter-subject normalisation

Given a manifold equipped with an affine connection, for example a Riemannian manifold with its Levi-Civita connection, one can define the parallel transport of a tangent vector v along a curve γ .

Definition 1.1 (Parallel vector field). *Let M be a smooth manifold and ∇ a connection on M . For any curve $\gamma : [a, b] \rightarrow M$ in M , a vector field X along γ is parallel if*

$$\nabla_{\dot{\gamma}(t)} X(t) = 0. \quad (1)$$

From the properties of ODEs, one can prove that given γ and v a tangent vector at $x = \gamma(0)$, there exists a unique parallel vector field X along γ such that $X(0) = v$. Intuitively, this ODE constrains the parallel vector field to keep its orientation w.r.t the velocity $\dot{\gamma}$ of the curve while moving along it. For any t in the domain of definition of γ , $X(t)$ is called the *parallel transport* of v at $y = \gamma(t)$ along γ , and written $\Pi_x^y v$.

For normalizing the momentum of a longitudinal deformation along an inter-subject geodesic diffeomorphism, parallel transport was first proposed in the LD-




Figure 2: Representation of the two ladder schemes, Schild’s (left) and pole ladder (right). The methods consist in iterating the construction of geodesic parallelograms to approximate the parallel transport of a vector v along a curve which is geodesic by arc.

DMM framework using Jacobi fields (Younes 2007; Qiu et al. 2008; Qiu et al. 2009). The method was used in (Cury et al. 2016) for the analysis of the thalamus in fronto-temporal dementia and it was improved by the Fanning Scheme (Louis et al. 2018), which is applied to shape analysis of brain structures (Louis et al. 2017). In latter works, the authors claim that the Jacobi Field approach is more precise and less expensive computationally than ladder methods. We have proved in Guigui and Pennec 2021 that this is not the case. Moreover, the proposed method implicitly assumes that the same metric is used for both longitudinal and inter-subject deformation. This is an important concern as the cross-sectional inter-subject and longitudinal intra-subject deformations have a fundamentally different nature.

In order to implement a parallel transport algorithm that remains consistent with the numerical scheme used to compute the geodesics, Lorenzi, Ayache, and Pennec 2011; Lorenzi and Pennec 2013 proposed to adapt Schild’s ladder to image registration with deformations parametrized by SVF. This method relies on iterating the construction of geodesic parallelograms to approximate the parallel transport of a vector v along a geodesic (see Figure 2a). Interestingly, the Schild’s ladder implementation appeared to be more stable in practice than the closed-form expression of the symmetric Cartan-Schouten parallel transport on geodesics. This was attributed to the inconsistency of the numerical schemes used for the computation of the transformation Jacobian in the implementation of this exact parallel transport formula.

Shortly after, it was realized that parallel transport along geodesics could exploit an additional symmetry by using the geodesic along which we want to transport as one of the diagonal of the geodesic parallelogram: this gave rise to the pole ladder scheme (Lorenzi and Pennec 2014). In this case the geodesic along which we are transporting is used as diagonal of the parallelograms (see Figure 2b). In this case the geodesic along which we are transporting is used as diagonal of the parallelograms (see Figure 2b). This greatly reduces the number of geodesics to compute ($-v_1$ does not even need to be computed on Figure 2b). Pole ladder combined with an efficient numerical scheme based on the Baker-Campbell-Hausdorff formula was found to




Figure 3: Statistics on diffeomorphisms with the Cartan-Schouten connection to model the normal component of the aging trajectory and the additional component specific to Alzheimer’s disease. Longitudinal geodesic deformations regressed in the sequence of images of each subject are parallel transported along the subject-to-reference deformation at baseline. In this common space, a linear model in the tangent space of diffeomorphisms estimates the mean trajectory for the two clinical conditions. Derived from original images and illustrations of Marco Lorenzi and Raphaël Sivera.

be more stable on simulated and real experiments than the other parallel transport schemes tested on Lie groups. This result and the higher symmetry led to conjecture that pole ladder could be a higher order scheme than Schild’s ladder.

The SVF framework was successfully applied in (Lorenzi et al. 2015) to distinguish pathological evolution from normal aging. An efficient computational pipeline is provided by (Hadj-Hamou et al. 2016) and used in (Sivera et al. 2020) to model Alzheimers’s disease and to analyze the effect of a potential treatment in a clinical trial (see illustration in Fig. 3).

1.4 Chapter organization

In this chapter, we address the longitudinal analysis of cardiac motion across patients and diseases. A key problem of this application is the large cyclic nature of cardiac deformations, along with large to very large inter-subject transformations. In order to caption the large cardiac motion, we chose to use the LDDMM framework because a right invariant metric is compatible with a Lagrangian and a Hamiltonian formulation of the mechanical equations, while the SVF framework is rather related to a Eulerian invariance.

The numerical complexity of the implementation of Jacobi fields and the algorithmic simplicity of ladder methods led us to choose the latter for inter-subject normalization. However, despite the practical success of ladder methods for longitudinal brain analyses, their numerical accuracy remained essentially unknown beyond the first order. This led Guigui and Pennec 2021 to conduct a careful numerical analysis of Schild’s and pole ladders that we summarize in section 2.1. To the best of our knowledge, ladder methods have not been used in the LDDMM framework, for which a BCH-type approximation was missing. We first give an overview of the LDDMM framework in section 2.2 before proposing a new second-order implementation of pole ladder on LDDMM deformations in section 2.3.

We turn in section 3 to the application of this methodology to the group-wise analysis of cardiac motion across diseases. For the cardiac motions that we analyse in section 3.2, relative volume differences between subjects appear to have an unexpected effect on the parallel transported trajectories. Since this undesirable effect cannot be attributed to the numerical accuracy of the transport scheme, it indicates that we have to revise the choice of the Riemannian metric used for the inter-subject comparison. We first investigate in section 3.3 an intuitive rescaling of the modulus of the momentum after the parallel transport to preserve on average the ejection fraction along the motion trajectory. Regression results over a population of subjects and patients show a surprisingly simple relationship between the scaling and the inter-subject volume ratio. However, modifying the transport equations in an ad-hoc fashion is not satisfactory from a theoretical point. Thus, we investigate in section 3.4 a more satisfactory strategy that decouples the volume change from the deformation directly within the LDDMM metric.

2 Parallel transport with ladder methods

Numerical methods have been proposed in geometric data processing for the parallel transport of vectors along curves in manifolds. The oldest algorithm is probably Schild’s ladder, a general method for the parallel transport along arbitrary curves, introduced in the theory of gravitation in Misner, Thorne, and Wheeler 1973 in the spirit of Schild’s geometric constructions. Ehlers, Pirani, and Schild 1972 is often cited as reference for Schild’s ladder but no mention of the scheme is made in this work. The method extends the infinitesimal transport through the construction of geodesic parallelograms (Figure 4). It is algorithmically interesting since it only requires the computation of geodesics (initial and boundary value problems) without requiring the knowledge of the second order structure of the space (connection or curvature tensors). Kheyfets, W. A. Miller, and Newton 2000 proved that the scheme realizes a first order approximation of the parallel transport for a symmetric connection. This makes sense since the skew-symmetric part of the connection, the torsion, does not impact the geodesic equation. Schild’s ladder is nowadays increasingly used in non-linear data processing and analysis to implement parallel transport in Riemannian manifolds. One can cite for instance Lorenzi, Ayache, and Pennec 2011 and the follow-up publications cited in introduction for the parallel transport of deformations in computational anatomy, or Hauberg, Lauze, and Pedersen 2013 for parallel transporting the covariance matrix in Kalman filtering.




Figure 4: Schild’s ladder procedure to parallel transports the vector v along the sampled curve with initial velocity w . **Top:** First rung of the ladder using an approximate geodesic parallelogram. **Bottom:** The method is iterated with rungs at each point sampled along the curve. Figures reproduced from Guigui and Penneec 2021

2.1 Numerical accuracy of Schild’s and pole ladders

Despite the use of ladder methods in applications, no results on their convergence were published before the work of Guigui and Penneec 2021 that we summarize here. We give Taylor approximations of the elementary constructions of Schild’s ladder and of the pole ladder with respect to the Riemann curvature of the underlying space. This allows to prove that these methods can be iterated to converge with quadratic speed, even when geodesics are approximated by numerical schemes.

Elementary construction of Schild’s ladder Recall that Exp denote the Riemannian exponential map, that maps a point and a tangent vector to the point reached at time $t = 1$ by the unique geodesic with those initial conditions. Log is its inverse defined locally, such that $\text{Exp}_x(\text{Log}_x(y)) = y$. See Penneec, Sommer, and Fletcher 2020, chapter 1 for more details. The construction to parallel transport $v \in T_x M$ along the geodesic γ with $\gamma(0) = x$ and $\dot{\gamma}(0) = w \in T_x M$ (such that $(v, w) \in U_x$) is given by the following steps (Figure 4):

1. Compute the geodesics from x with initial velocities v and w until time $s = t = 1$ to obtain x_v and x_w . These are the sides of the parallelogram.
2. Compute the geodesic between x_v and x_w and the midpoint m of this geodesic:

$$m = \text{Exp}_{x_v} \left(\frac{1}{2} \text{Log}_{x_v}(x_w) \right).$$

This is the first diagonal of the parallelogram.

3. Compute the geodesic between x and m , let $a \in T_x M$ be its initial velocity. Extend it beyond m for the same length as between x and m to obtain z :

$$a = \text{Log}_x(m); \quad z = \text{Exp}_x(2a) = x_{2a}.$$

This is the second diagonal of the parallelogram.

4. Compute the geodesic between x_w and z . Its initial velocity u^w is an approximation of the parallel transport of v along the geodesic from x to x_w , i.e.

$$u^w = \text{Log}_{x_w}(x_{2a}).$$

Assuming that there exists a convex neighborhood that contains the entire parallelogram, all the above operations are well defined. In the literature, this construction is then iterated along γ without further precision.

Taylor expansion We can now reformulate this elementary construction in terms of successive applications of the double exponential and the neighboring logarithm maps, given in Penneec 2019. The computations are detailed in the appendix of Guigui and Penneec 2021, and we only report here the result at third order, meaning that all the terms of the form $\nabla.R(\cdot, \cdot)$ are summarized in the term $O(4)$. We first obtain a generalized midpoint rule:

$$2a = w + v + \frac{1}{6}R(v, w)(w - v) + O(4) \quad (2)$$

We notice that this expression is symmetric in v and w , as expected. Furthermore, the deviation from the Euclidean mean of v, w (the parallelogram law) is explicitly due to the curvature. Accounting for this correction term is a key ingredient to reach a quadratic speed of convergence.

Now, by propagating this midpoint rule to compute the error e_x made by this construction to parallel transport v , $e_x = \Pi_x^{x_w}v - u^w$, we obtain a third order approximation of the Schild's ladder construction:

Theorem 2.1. *Let (M, g) be a finite dimensional Riemannian manifold. Let $x \in M$ and $v, w \in T_x M$ sufficiently small. Then the output u of one step of Schild's ladder parallel transported back to x is given by*

$$u = v + \frac{1}{2}R(w, v)v + O(4) \quad (3)$$

This theorem shows that Schild's ladder is in fact a second-order approximation of parallel transport. Furthermore, this shows that splitting the main geodesic into segments of length $\frac{1}{n}$ and simply iterating this construction n times will in fact sum n error terms of the form $R(\frac{w_i}{n}, v_i)v_i$, hence by linearity of R , the error won't necessarily vanish as $n \rightarrow \infty$. To ensure convergence, it is necessary to also scale v in each parallelogram, as detailed in the next paragraph.

Numerical scheme and convergence With the previous notations, let us define $\text{schild}(x, w, v) = u^w \in T_{x_w}M$. We now divide the geodesic γ into segments of length $\frac{1}{n}$ for some $n \in \mathbb{N}^*$ large enough so that the previous Taylor expansions can be applied to $\frac{w}{n}$ and $\frac{v}{n}$. As mentioned before, v needs to be scaled as w . In fact let $\alpha \geq 1$ and consider the sequence defined by (see Figure 4)

$$\begin{aligned} v_0 &= v \\ v_{i+1} &= n^\alpha \cdot \text{schild}\left(x_i, \frac{w_i}{n}, \frac{v_i}{n^\alpha}\right), \end{aligned} \quad (4)$$




Figure 5: **Top:** Elementary construction of the pole ladder with the previous notations (left), and new notations (right), in a normal coordinate system at m . **Bottom:** Iterated scheme. Figures reproduced from Guigui and Penneec 2021.

where $x_i = \gamma(\frac{i}{n}) = \text{Exp}_x(\frac{i}{n}w)$, $w_i = n \text{Log}_{x_i}(x_{i+1}) = \Pi_x^{x_i}w$. We now establish the following result, which ensures convergence of Schild's ladder to the parallel transport of v along γ at order at most two.

Theorem 2.2. *Let $(x_i, v_i, w_i)_{(i \leq n)}$ be the sequence defined as above. Then $\exists \tau > 0, \exists \beta > 0, \exists N \in \mathbb{N}, \forall n > N$,*

$$\|v_n - \Pi_x^{x_n}v\| \leq \frac{\tau}{n^\alpha} + \frac{\beta}{n^2}.$$

Moreover, τ is bounded by a bound on the sectional curvature, and β by a bound on the covariant derivative of the curvature tensor.

The same result can be obtained for pole ladder, and in the case where geodesics are obtained by numerical integration of the geodesic equation.

Pole ladder In this case, the main geodesic is used as diagonal of the constructed parallelograms (see Figure 5). This allows to reduce the number of geodesics that need to be computed. Similarly, a Taylor approximation of the construction gives

$$u = v + \frac{1}{12}((\nabla_w R)(w, v)(5v - w) + (\nabla_v R)(w, v)(2v - w)) + O(5). \quad (5)$$

And this can be propagated to show the following bound

$$\|v_n - \Pi_m^{m_n}v\| \leq \frac{\beta}{n^2}, \quad (6)$$

that ensures convergence to the parallel transport with a quadratic speed.

Infinitesimal schemes When geodesics are not available in closed form, we replace the Exp map by a fourth-order numerical scheme (e.g. Runge-Kutta (RK)), and the Log is obtained by gradient descent over the initial condition of Exp. It turns out that only one step of the numerical scheme is sufficient to ensure convergence, and keeps the computational complexity reasonable. As only one step of the

integration schemes is performed, we are no longer computing geodesic parallelograms, but infinitesimal ones, and thus refer to this variant as *infinitesimal scheme*. As in the previous case, we can show

$$\|\Pi_x^{x_n} v - v_n\| = O\left(\frac{1}{n^2}\right).$$

This allows to use this scheme in the SVF and LDDMM frameworks. For the later, we first give an overview of the framework in the next section before presenting our parallel transport algorithm.

2.2 A short overview of the LDDMM framework

The LDDMM framework encompasses both algorithms for shape matching (also referred to as registration) and a Riemannian geometric structure on the space of deformations, that projects to the space of shapes. The former allows computing shape descriptors, and to parameterize diffeomorphisms. The latter, provides a distance to compare deformations, and an associated notion of parallelism to transport them. We summarize here the formulation of Durrleman, Allasonnière, and Joshi 2013; Durrleman et al. 2014 and their implementation in Bône et al. 2018, focusing on the case of landmarks. For a thorough treatment of the topic, we refer to Younes 2019.

In this chapter, we restrict to shapes defined as $d - 1$ -dimensional submanifolds of \mathbb{R}^d , $d = 1, 2$, with the same given topology, and approximated by triangulated meshes, defined by a set of points in \mathbb{R}^d , called landmarks, and connectivity between those. The set of shapes is denoted \mathcal{S} . The model of computational anatomy stipulates that such a shape is fully described by a diffeomorphism $\varphi \in \text{Diff}(\mathbb{R}^d)$, where the interaction between shapes and diffeomorphisms is encoded by the action of $\text{Diff}(\mathbb{R}^d)$ on \mathcal{S} . This turns the shape space into a homogeneous space with group $\text{Diff}(\mathbb{R}^d)$. The isotropy subgroup of a given shape is the set of re-parameterization of this shape, i.e. diffeomorphisms that leave the shape unchanged, but move particles along the shape. These diffeomorphisms are infinite dimensional nuisance parameters in a statistical analysis, and one feature of LDDMM is to provide methods that are invariant to such nuisance (M. Miller, Trounev, and Younes 2015).

A natural and efficient computational construction of diffeomorphisms is obtained by flows associated to ordinary differential equations (ODEs) $\partial\phi_t(\cdot) = v_t[\phi_t(\cdot)]$, with the initial condition $\phi_0 = Id$. The time-dependent vector field v_t can be interpreted as the instantaneous speed of the points during deformation, and must verify certain regularity conditions to ensure that solutions to the ODE are indeed diffeomorphisms. An efficient way to enforce these conditions is to restrict to vector fields obtained by the convolution of a number N_c of momentum vectors carried by distinct control points:

$$v_t(x) = \sum_{k=1}^{N_c} K(x, c_k^{(t)}) \mu_K^{(t)}, \quad (7)$$

where K is a kernel function, e.g. the Gaussian kernel: $K(x, y) = \exp\left(-\frac{\|x-y\|^2}{\sigma^2}\right)$. The effect of the choice of kernel is studied in Micheli and Glaunès 2014, but we restrict to the Gaussian kernel in this study. The (closure of the) set of such vector

fields forms a reproducing kernel Hilbert space H_K , with the associated norm

$$\|v\|_K^2 = \sum_{i,j} K(c_i, c_j) \mu_i^T \mu_j.$$

We write Diff_K the subgroup of diffeomorphisms obtained this way and use matrix notation with a $dN_c \times dN_c$ block matrix $\mathbf{K}(c, c') = (K(c_i, c'_j) I_d)_{ij}$ and flat vectors of size dN_c for landmarks sets, velocities and momentum, so that $v(x) = \mathbf{K}(c, x)\mu$ and $\|v\|_K^2 = \mu^T \mathbf{K}(c, c)\mu$. The total cost, or energy of the deformation can be defined as $\int_0^1 \|v_t\|_K^2 dt$.

It can be shown that the momentum vectors that minimize this energy, considering $c_k^{(0)}, c_k^{(1)}$, $k = 1 \dots N_c$ fixed, together with the equation driving the motion of the control points, follow a Hamiltonian system of ODEs:

$$\begin{cases} \dot{c}^{(t)} &= \mathbf{K}(c^{(t)}, c^{(t)})\mu^{(t)} \\ \dot{\mu}^{(t)} &= -\nabla_1 \mathbf{K}(c^{(t)}, c^{(t)})\mu^{(t)T} \mu^{(t)} \end{cases} \quad (8)$$

A diffeomorphism ϕ_1 is thus uniquely parameterized by the initial conditions $c_k^{(0)}, \mu_k^{(0)}$, $k = 1 \dots N_c$, and a shape registration criterion between a template \bar{q} and target q can be defined as

$$C(c, \mu) = \|q - \phi_1^{c, \mu}(\bar{q})\|_2^2 + \alpha^2 \|v_0^{c, \mu}\|_K^2. \quad (9)$$

where α is a regularization parameter that penalises large deformations. Minimizing C therefore amounts to finding the transformation that best deforms \bar{q} to match q . The gradient of C can be computed through automatic differentiation, to perform gradient descent. Note that in (9), the L_2 norm between landmark positions is used to evaluate the data-attachment term. When corresponding landmarks are not available, fidelity metrics relying on currents, varifolds or normal cycles have been derived and allow to compute metrics between curves or surfaces with different parametrizations. These are reviewed in Charon et al. 2020.

Two examples of registration solutions are shown, the first is on simple parametric shapes, a circle and an ellipse (Figure 6). In this case the full trajectory of each landmark is represented (with a color scheme from blue to red). These differ from a linear interpolation especially where the curvature of the shape increases. The second example is on real RV data (Figure 7). The distribution of momentum vectors on the shape is not intuitive, and it is easier to inspect the velocity field evaluated at the control points. The matching is very efficient on these meshes, thanks to the quality of the correspondences between landmarks and to the smoothness of the shapes.

The optimal value of C defines a distance between ϕ_1 and the identity. In fact this distance derives from a right-invariant Riemannian metric on the group Diff_K and the path $t \mapsto \phi_t$ is a minimizing geodesic for this metric. Integrating the system (8) computes the exponential map of this metric, and is often called *shooting* in this section, while minimizing (9) approximates the logarithm, and is called *registering*. By considering the action of diffeomorphisms on shapes, this distance projects to a distance between the shapes q and \bar{q} , and we have

$$d_K(q, \bar{q})^2 = \inf \left\{ \frac{1}{2} \int_0^1 \dot{q}_t \mathbf{K}(c, c)^{-1} \dot{q}_t dt \mid q_0 = \bar{q}, q_1 = q \right\}$$




Figure 6: Example of registration and optimal path between a circle (source) and an ellipse (target) with 60 landmarks, 15 control points (green) and corresponding momentum vectors (cyan) and kernel width 1. The progression from blue to red shows the time trajectories of the landmarks under the minimizing geodesic deformation.

This in fact turns the space of shapes into a homogeneous space, with invariant metric given by the group.


2.3 Ladder methods with LDDMM

Along with a distance, the Riemannian metric provides a notion of parallel transport thanks to its Levi-Civita connection. A Hamiltonian formulation of parallel transport may be given by an ODE that allows to transport a set of momentum vectors along a path of diffeomorphisms ϕ_t (Younes 2019, section 13.3.3). Previous implementations of parallel transport of deformations in the context of LDDMM are however based on Jacobi fields (Younes 2007; Cury et al. 2016), and in particular on the *fanning scheme* (Louis et al. 2018).

Ladder method in contrast, were proposed in the SVF framework. Adapting them to LDDMM is difficult because of the absence of an equivalent of the BCH formula to implement each rung of the ladder. We propose here a new implementation of pole ladder in the context of LDDMM that leverages the lower dimensional parametrization and the Hamiltonian formulation, and is grounded by the results of the previous section. Indeed, the geodesics are not known in closed-form and we are in the case of infinitesimal ladder schemes, requiring the use of fourth-order integration schemes. Our implementation is largely inspired by that of the fanning scheme and leverages the open-source software Deformetrica (Bône et al. 2018).

Throughout this section, we consider a reference shape q_A , and wish to transport the deformation (c, μ) obtained by registration of a shape q_1 on another shape q_0 to q_A , that is, along the geodesic between q_0 and q_A , parametrized by (c_A, ω) and also obtained by registration. We suppose that registration outputs a set of control points such that $\mathbf{K}(c, c)$ is a symmetric positive definite matrix.

Recall from the homogeneous space structure that the tangent space at q_0 is identified with the horizontal subspace of the Lie algebra of Diff_K , which is here the reproducing Kernel Hilbert space H_K that contains all vector fields defined by the



(a) Systolic deformation with momenta (b) Systolic deformation with velocity field

Figure 7: Example of the registration between ED (blue wireframe) and ES (red points) for a control right-ventricle. There are 938 landmarks, 34 control points with kernel width 10mm (the overall ventricle at ED is about 55mm high). Note how momentum vectors are noisy due to strong interactions between control points (left). The velocity field at the control points offers a better summary of the deformation (right).




Figure 8: Representation of the pole ladder for diffeomorphisms with an odd number of steps n . The exponential maps are computed with a RK4 scheme, and the log by gradient descent.

convolution (7). It defines a one-to-one correspondence between the tangent space at q_0 , and its co-tangent space, meaning that we can transport the momentum μ and flow the points c instead of transporting a full vector field v . Similarly, we don't need to flow the full shape q_0 along the ladder, but only the control points.

We describe the pole ladder in the case where the momentum to transport μ and the momentum along which we transport ω are carried by the same initial control points $c = c_A$. If it is not the case, μ is projected to c_A by solving a linear system with a Cholesky decomposition of $\mathbf{K}(c_A, c_A)$: $\mu' = \mathbf{K}(c_A, c_A)^{-1} \mathbf{K}(c, c_A) \mu$.

Define $\text{rk} : TS^* \times \mathbb{R}_+ \rightarrow TS^*$ the map that performs one step of the fourth-order Runge-Kutta numerical scheme on the Hamiltonian system (8), and rk_1 the projection on the control points only, and rk_x^{-1} its approximate inverse by gradient descent, such that $\|\text{rk}_1(\text{rk}_x^{-1}(y), h) - y\| \leq h^5$. Then, pole ladder is performed by the following algorithm.

Choose $n \in \mathbb{N}$ and $1 \leq \alpha \leq 2$, and divide $[0, 1]$ into n intervals of length $h = \frac{1}{n}$. Compute the first midpoint $c_1, \omega_1 = \text{rk}(c_A, \omega, \frac{h}{2})$ along the main geodesic, and the




Figure 9: Parallel transport of the evolution from a circle to an ellipse along the deformation to a smaller circle.

first rung $q_0 = \text{rk}_1(\mu, \frac{1}{n\alpha})$. Then iterate for $i = 1 \dots n$ (see Figure 8)

1. Compute the momentum $\alpha_i = n \text{rk}_{c_i}^{-1}(q_{i-1})$;
2. Compute the flow of its inverse $q_i = \text{rk}_1(c_i, -\alpha_i, h)$;
3. Compute the next midpoint $c_{i+1}, \omega_{i+1} = \text{rk}(c_i, \omega_i, h)$ except at the last step where $h \leftarrow \frac{h}{2}$ to compute \tilde{c} ;

Return $\tilde{\mu} = n^\alpha (-1)^n \text{rk}_{\tilde{c}}^{-1}(q_n)$.

Two examples of parallel transport solutions are shown, the first on simple parametric shapes, the evolution from a circle to an ellipse is transported to a smaller circle (Figure 9). We retrieve a smaller ellipse as expected. The second example on real right ventricle (RV) mesh of the heart (Figure 10), shows the ED-to-ES deformation transported to a third reference shape, the atlas. The obtained reconstruction is a personalized estimate of an ES frame for the atlas, specific to the considered patient. The transported frame is much more acceptable than those of Figure 1.

Validation Recall that parallel transport is an isometry, so the norm of the transported deformation must equal that of the initial deformation. We use this property as a first step to validate our implementation on a population of simulated 2d-shapes. These are generated by shooting from a circle with random Gaussian momentum vectors. This defines the source shape. Then an affinity is applied to the source shape to define its evolution (i.e. scaling with different coefficients along each axis). This deformation is then estimated by registration and transported to the template. The reconstruction obtained by deforming the circle with the transported deformation (yellow shape on the figure) resembles an ellipse, as expected by the application of the affinity.

For $n = 10$ rungs, we obtain a root mean squared error (RMSE) of 8.3×10^{-3} and less than $1.3 \pm 0.9\%$ relative error (in absolute value) for 100 samples. We also apply the affinity to the template and register the template to the result to compute the expected momentum after transport. The deviation of the transported momentum to this expected value is measured with the kernel norm of the difference and averaged over the samples. We obtain an error of 0.34 while the fanning scheme

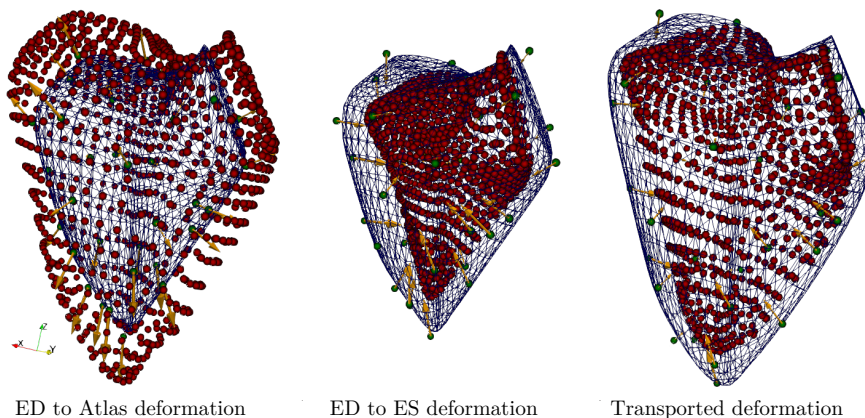


Figure 10: Example ED-to-ES deformation transported along the ED-to-atlas path. The source frames are represented by a blue wireframe, the final frames by red landmarks, the control points by green landmarks and the initial velocities at the control points by orange arrows. The transported frame is much more acceptable than those of Figure 1.

achieves 0.33. The two implementations are therefore very similar on this set of shapes. We however find cases where the fanning scheme does not behave as well as the pole ladder for more complex deformations on the heart data (Figure 13 below).

3 Application to cardiac motion modelling

Spatio-temporal shape analysis is of growing importance in the study of cardiac diseases. In particular, the assessment of cardiac function requires the measurement and analysis of cardiac motion beyond scalar indicators such as volume, ejection fraction, pressure-volume loops or area strain. We focus on the case of the right ventricle (RV), that is of particular interest as it has been shown to have a large capacity to adapt to overload by remodelling (Sanz et al. 2019), raising the issue of disentangling the deformation from the initial anatomy.

3.1 The Right Ventricle and its diseases

The right ventricle receives the blood from the right-atrium through the tricuspid valve and ejects it to the lungs through the pulmonary artery. Although neglected in the assessment of normal blood circulation of adult patients for decades, its relevance to determine symptoms and outcomes is now well established in several pathologies. Three such pathologies are studied in our database.

Pulmonary hypertension (PH) is a severe condition characterized by increased blood pressure in the pulmonary arteries, appearing between the ages of thirty and forty. It is associated with symptoms and premature deaths (Moceri et al. 2018) and the most common source of pressure overload in the right-ventricle, resulting in hypertrophy, flattening of the interventricular septum, progressive dilation and dysfunction (Sanz et al. 2019).

Atrial septal defect (ASD) is a congenital heart defect in which blood flows between the left and right atrium. This may lead to lower oxygen levels in the blood

that supplies the other organs, resulting in cyanosis. It affects the RV with volume overload, marked by dilation and hypertrophy and leftward septal shift, but it may be tolerated for years, as many studies report no affectation of the function of the RV under volume overload (Moceri et al. 2020; Sanz et al. 2019).

Tetralogy of Fallot (ToF) is another congenital heart defect characterized by pulmonary stenosis (i.e. narrowing of the pulmonary valve) and an overriding aorta, allowing blood flows between the left and right ventricles, and resulting in hypertrophy of the RV, reduced oxygen level and cyanosis. This is the unique case where the RV is under both volume and pressure overload. Untreated, ToF patients rarely survive to adulthood. It is treated by surgical repair during the first year of life, greatly improving survival, but post-surgery defects such as pulmonary regurgitation may occur, requiring other operations.

We use a database of 3D meshes extracted from 314 echocardiographic sequences from patients examined at the CHU of Nice (see Moceri et al. 2018; Moceri et al. 2020 for details). The meshes were extracted with a commercial software (4D RV Function 2.0, TomTec Imaging Systems, GmbH, DE) with point-to-point correspondences across time and patients. These are formed of 938 points and 1872 triangles. All the shapes were realigned with a subject-specific rigid-body deformation. An atlas was computed from the end diastolic meshes of the control group, after rigid-body alignment. We warmly thank Pamela Moceri and Nicolas Duchateau for collecting and curating the data, and for the helpful interactions that we had during this project.

Two main measures used by the clinicians to evaluate the cardiac function will be used in the sequel to evaluate our modelling pipeline. *Ejection fraction (EF)* is the relative volume change between end-diastolic (ED) time and end-systolic (ES) time, i.e. during one contraction. It represents the volume of blood pumped by the heart to the body circulation, and is used in clinical routines to evaluate heart failures. *Area strain (AS)* is the relative area change of each cell of the mesh between ED time and ES time. As it depends on the quality of the mesh, it is usually filtered by computing the mean at each vertex over the neighboring cells, and visualized as a colormap on the mesh, or averaged by regions of the ventricle. It represents the amount of stretching local tissues undergo during the deformation, and is a common descriptor of cardiac motion (Kleijn et al. 2011; Di Folco et al. 2019; Moceri et al. 2020).

3.2 Motion normalization with parallel transport

For each subject, the motion of the right ventricle during the systolic phase of the cardiac cycle is encoded by the LDDMM deformation at a discrete set of sampling times between ED and ES. Each time-point is treated as an independent observation for its reorientation by parallel transport into the atlas reference frame. The framework is summarized in Figure 11. For each observation time t_i , the subject’s frame at time t_i is registered to this subject’s ED frame, transported along the path between ED and the atlas, and reconstructed by shooting with the transported momenta.

We reproduce on Figure 12 the reconstructions of the ES frame for a patient of each disease group and of the control group. These are obtained by the pole ladder algorithm (right) as well as the fanning scheme (middle). On these examples the two methods produce very similar results, which validates our implementation.




Figure 11: Framework using registration and parallel transport to normalize cardiac deformations.

Moreover, the obtained reconstructions are smooth (compare with Figure 1) and seem to adapt the deformation pattern of each patient to the atlas.

In most cases, little difference can be observed qualitatively. We find however a few examples where the fanning scheme did not achieve to transport the deformation, and some points diverge or collapse in the reconstructions. This occurs when the original shapes are significantly larger than the atlas, and with strong spherification of the apex, causing large deformations as for the ASD patient of Figure 13.

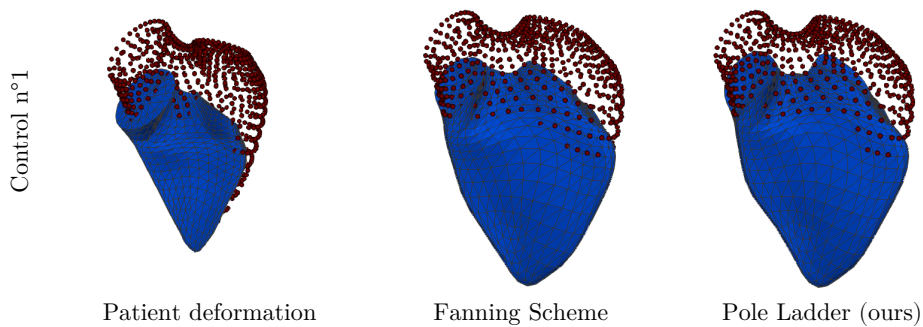


Figure 12: Examples of reconstructions of the ES frame (blue meshes) after parallel transport of the ED-to-ES deformation (point cloud to mesh) along the deformation from ED to atlas. Parallel transport is computed by the fanning scheme (middle) and our implementation of pole ladder (right). In most cases, little difference can be observed visually.




Figure 13: Examples of reconstruction where the fanning scheme did not achieve to transport the deformation, while pole ladder did, although it is not realistic.

Interaction between shape and deformations: a scale problem Niethammer and Vialard 2013 showed that parallel transport in LDDMM does not conserve global properties such as scale or volume changes. This is visible in our previous experiment (Figure 9), where the anisotropy of the original evolution (from circle to ellipse) is 1.4 whereas the one obtained after parallel transport is 2.3. In the case of cardiac deformations, the magnitude of the temporal deformation is comparable to that of the subject’s reference to atlas deformation, and substantial volume changes are observed. Thus the lack of scale-invariance is crucial.

Furthermore, there is a significant correlation between the magnitude of the systolic deformation and the ED volume ($\rho = 0.42$ in our dataset). Additionally, there are significant differences of ED volumes by disease groups (Figure 15a). As the parallel transport is isometric, this deformation may be too large for the atlas. Examples of the obtained ES frames are shown on Figure 14 for patients whose RV volume is greater than that of the atlas, resulting in non-realistic ES frames.

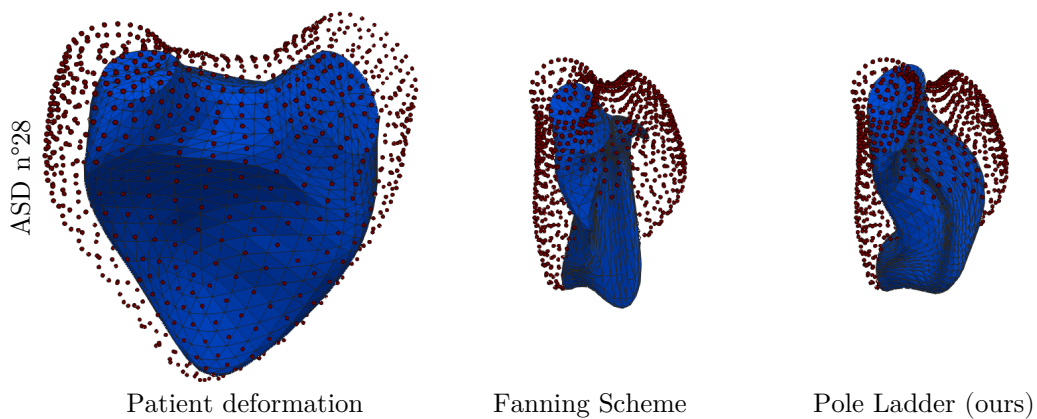


Figure 14: Examples of reconstructions of the ES frame (blue meshes) after parallel transport of the ED to ES deformation (point cloud to mesh) along the deformation from ED to atlas. The subjects belong to the ASD group, hence show large volume changes that result in unrealistic ES frames.




Figure 15: End-diastolic volume by disease group (left), original ejection fraction (middle) and its modification by parallel transport (right).

To quantify this phenomenon, we measure the EF and the AS before and after transport and report the RMSE in Table 1. We observe a large relative RMSE EF, representing 31% of the mean value, and equal to the standard deviation of the data. Furthermore, plotting the absolute value of the

EF error by disease group, we retrieve a strong correlation between the error and the EF (Figure 15c). This effect is quite undesirable for a normalization procedure!

In conclusion, parallel transport was able to reorient the deformation to a different frame, but was not sufficient to properly normalize it in the presence of large volume differences. Since we control the numerical accuracy of the parallel transport method thanks to 2.1, the error has to be attributed to the choice of the transport method or to the choice of the Riemannian metric used on deformations. In order to achieve proper normalization for a broader range of volume changes and to preserve the relative volume changes, we investigate in Subsection 3.3 a first intuitive ad-hoc strategy where we modulate the amplitude of the parallel transported vector. However, this amounts to state that the transport method and the Riemannian metric are not anymore consistent together. To remove this inconsistency with the general geometric statistics framework, we then investigate in Section 3.4 a modification of the metric on diffeomorphisms that decouples the volume change from the deformation.

3.3 An intuitive rescaling of LDDMM parallel transport

In Guigui et al. 2021, we hypothesises that the amplitude of the ED to atlas deformation, i.e. the one along which we transport, acts on the transported deformation by scaling to construct a normalization strategy invariant to initial volume changes.

Hypothesis We now propose to introduce a rescaling step after the parallel transport in our framework of Figure 11. More precisely, instead of using $R_t = \text{Exp}_A(\Pi_{ED}^A \text{Log}_{ED}(S_t))$ as the subject specific reconstruction of time-frame S_t , we introduce a parameter $\lambda > 0$ and use

$$R_t(\lambda) = \text{Exp}_A(\lambda \Pi_{ED}^A \text{Log}_{ED}(S_t)).$$

	Original Values	RMSE PT
AS	-0.24 ± 0.08	0.18
EF	0.42 ± 0.13	0.13

Table 1: Validation of parallel transport with EF and AS.

where Π_{ED}^A is the parallel transport map along the geodesic that joins the ED frame to the atlas. Based on our observation of the results of parallel transport (Figures 12 and 14) and of the relationship between the ejection fraction error and the disease group (Figure 15c), we hypothesise that this rescaling should be patient-specific and depend on the volume of the ED frame. Furthermore, rescaling should not be necessary if the volume of the ED frame matches that of the atlas. The relevant quantity is thus the ratio $\text{Vol}_{ED} / \text{Vol}_A$.

Criterion and Estimation of λ Given a patient’s sequence, we now need a criterion to find λ . As discussed in the introduction, a clinically relevant quantity that should be conserved is the ejection fraction, as it is a straightforward scalar indicator of cardiac function. A perfect preservation of the ejection fraction between each time point and the ED frame constraints the volume of each reconstruction:

$$\text{Vol}_{new}(t) = \frac{\text{Vol}_A}{\text{Vol}_{ED}} \text{Vol}_{S_t}.$$

A straightforward loss function for λ is therefore

$$\mathcal{L}(\lambda) = \sum_{i=1}^k (\text{Vol}(R_{t_i}(\lambda)) - \text{Vol}_{new}(t_i))^2.$$

For each patient, we then minimize this loss function with respect to λ to find the patient-specific scaling parameter. We solve this problem by gradient descent, where the gradient is computed by automatic differentiation. The parameter λ is regularized to be close to 1 to avoid poor solutions.

Results To validate the method, we first check that the ejection fraction is effectively conserved. The errors on ejection fraction and area strain are reported in Table 2. We indeed observe a three-fold decrease of the RMSE on the ejection fraction. Additionally, the distribution of this error is no longer related to the disease group (Figure 16). Visually, the reconstructions obtained by the scaled parallel transport are more realistic (Figure 17). Interestingly, it is possible to obtain a low error on the EF after the scaled transport, but this does not preserve the area strain (AS). This shows that although these quantities are related, they carry different information and the AS depends on the initial shape, itself related to the pathology.




Figure 16: Alteration of the ejection fraction per disease group. The scaled parallel transport achieves an error that is not related to the disease group.

Relationship between λ and Vol_{ED} As expected, the scaling coefficient is closely related to the ED volume. We use a linear regression to identify a linear relation between $\log(\lambda)$ and $\log(\text{Vol}_A / \text{Vol}_{ED})$. This is displayed on Figure 18. The

	Original Values	RMSE PT	RMSE SPT
AS	-0.24 ± 0.08	0.18	0.13
EF	0.42 ± 0.13	0.13	0.04

Table 2: Validation of scaled parallel transport (SPT) with EF and AS.

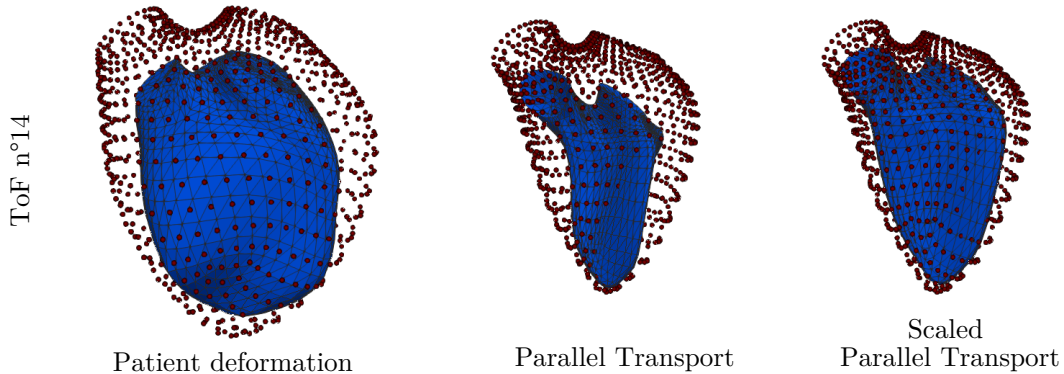


Figure 17: Examples of reconstructions of the ES frame (blue meshes) after parallel transport of the ED-to-ES deformation (point cloud to mesh) along ED-to-atlas deformation. The large volume changes between subjects and the atlas result in unrealistic ES frames (middle column). This problem is well addressed by the scaling strategy (right column).

relationship seems highly significant as we measure a coefficient of determination $R^2 = 0.92$. The measured slope is $\alpha = 0.64$ and corresponds to the relationship

$$\lambda = \beta \left(\frac{\text{Vol}_A}{\text{Vol}_{ED}} \right)^\alpha. \quad (10)$$

where β is the exponential of the intercept, but is found close to 1, with $\beta = 1.08$

In order to better understand this relationship, we reproduce the experiment on 2d simulated shapes, with the same experimental set-up as in the validation experiment. It is striking that we also find a highly significant relationship with $R^2 = 0.99$ and almost the same slope $\alpha = 0.67$ and $\beta = 1.02$. The values measured for α are close to $2/3$ and surprisingly do not seem to depend on the dimension of the ambient space.

Our experiments thus give strong evidence in favour of relation (10), and raise many questions. First is α a constant, or does it depend on the data, the dimension, the kernel, or the deformation model? From our results it seems that it does not depend on the data or the dimension. Performing the same experiments with different kernels could bring further information, this is ongoing work. Similarly, experiments in the SVF framework could be performed to test whether this effect is specific to LDDMM or more general. These experiments could also clarify whether β significantly differs from one.

Furthermore, this effect suggests that all the experiments that have been performed with parallel transport should be revised with a scaling step. Most of these




Figure 18: Scaling Parameter λ with respect to the ED volume.

in the medical imaging field concern the evolution of brain images (Cury et al. 2016; Lorenzi et al. 2015). However, longitudinal deformations of the brain are very different from that of the right ventricle, and we don't expect the scaling step to be necessary. Indeed, brain deformations are characterized by very localized strong Jacobian determinant, but very little overall volume changes. A more relevant type of data to reproduce this effect could be to evaluate tumor evolution during treatment, or the evolution of more plastic organs such as the liver, whose shape is affected by all the liver diseases, whether viral, metabolic or due to alcohol.

Finally, a last hypothesis is that by scaling the parallel transport of the LDDMM metric, we are approximating the parallel transport of a direct product metric with a volume preserving factor and a metric on the volume changes as will be introduced in the next section. This relation would translate in setting the metric on the volume part.

3.4 Changing the metric to preserve relative volume changes

Although modifying the parallel transport equations works astonishingly well from the practical point of view, it is not satisfactory from a theoretical point of view since it decorrelates the transport from the metric used on the deformations. We investigate in this section a more satisfactory strategy that decouples the volume change from the deformation directly within the LDMMM metric. This strategy is based on the scale invariant metric proposed by Niethammer and Vialard 2013 and it was implemented thanks to the precious advice of François-Xavier Vialard. More generally, they show that a Riemannian metric on the space of shapes preserves a global non-degenerate function f (such as volume) if and only if it can be decomposed into a product metric on $\text{Im}(f)$ and $f^{-1}(0)$ (see Niethammer and Vialard 2013, Theorem 41).

Model We apply their result as follows. Restricting to the space of shapes that are embeddings of the circle in \mathbb{R}^2 or of the sphere in \mathbb{R}^3 , let $M = \text{Emb}(S^d, \mathbb{R}^{d+1})$ be the set of shapes, endowed with a metric g . As we are interested in relative volume changes, we consider the function $f = \log\left(\frac{\text{Vol}}{\text{Vol}_A}\right) : M \rightarrow \mathbb{R}$ (so that $df = \frac{d\text{Vol}}{\text{Vol}}$),

and let $M_0 = f^{-1}(0)$, the space of shapes whose volume equals the volume of the atlas Vol_A . Then, by Niethammer and Vialard 2013, Theorem 41, (M, g) can be decomposed into a direct product of Riemannian manifolds $(M, g) = (\mathbb{R}, dt^2) \times (M_0, g_0)$ where g_0 is a Riemannian metric on the submanifold M_0 . This metric can be constructed by choosing a projection $\pi : M \rightarrow M_0$ and the restriction of g to M_0 . However, there is no canonical projection and two schemes are proposed in Niethammer and Vialard 2013:

- by gradient flow: we follow the flow of $(d \text{Vol})^\#$ where $\#$ depends on the choice of metric on M , in this case the LDDMM metric.
- by scaling: we centre the shape around 0 and divide all the landmarks by $\text{Vol}_q^{1/3}$. This choice depends on the centre (barycenter) of the shape which may be unnatural.

In both cases, the framework is slightly modified by applying this projection first to all the time frames, so that their volume matches that of the atlas. Then the previous framework of Figure 11 is applied with volume preserving geodesics and parallel transport. This corresponds to the parallel transport on the factor M_0 of $M = \mathbb{R} \times M_0$. Finally, the (Euclidean) transport on \mathbb{R} corresponds to applying the vector $\log\left(\frac{\text{Vol}_{S_t}}{\text{Vol}_{ED}}\right)$ to 0, so that the new volume is $\text{Vol}_{new}(t) = \frac{\text{Vol}_A}{\text{Vol}_{ED}} \text{Vol}_{S_t}$ as in the previous section, and the ejection fraction is preserved by construction. This volume is obtained either by scaling or by gradient flow, consistently with the projection.

Implementation We now give more details on the computation of geodesics and our implementation in the LDDMM framework of Section 2.2. Recall that tangent spaces are described by the Hilbert space H_K of vector fields obtained by convolution. As we restrict to $M_0 = f^{-1}(0)$, a vertical tangent subspace at any shape q is defined as the set of vector fields of H_K that are volume preserving, i.e.

$$\ker(df_q) = \{v \in H_K \mid d \text{Vol}_q((v(q))) = 0\},$$

To avoid any confusion, we distinguish the linear form $d \text{Vol}_q$ from its representation in the canonical basis ∂Vol_q , such that $d \text{Vol}_q(v(q)) = v(q)^\top \partial \text{Vol}_q$. The orthogonal projection on V_q (which depends on the LDDMM metric) can be computed by

$$\begin{aligned} \pi_q : T_q M &\longrightarrow V_q \\ v &\longmapsto v - d \text{Vol}_q(v) \vec{n} \end{aligned} \quad (11)$$

where $\vec{n} = \frac{(d \text{Vol}_q)^\#}{\|(d \text{Vol}_q)^\#\|^2}$. Note that the map $\#$ associates a tangent vector to the linear form $d \text{Vol}$ such that $d \text{Vol}(v) = \langle (d \text{Vol}_q)^\#, v \rangle_K$, so that $(d \text{Vol}_q)^\# = \mathbf{K}(\cdot, q) \partial \text{Vol}_q$. Define also the dual $\pi_q^* : T_q^* M \rightarrow V_q^*$ to π_q by the relation $\mu(\pi_q(v)) = \pi_q^*(\mu)(v)$, i.e. in matrix notations $\mu^\top \pi_q(v) = v^\top \pi_q^*(\mu)$. From (11), we obtain

$$\begin{aligned} \pi_q^* : T_q^* M &\longrightarrow V_q^* \\ (c, \mu) &\longmapsto \left(c, \mu - \frac{\mu^\top \mathbf{K}(c, q) \partial \text{Vol}_q}{\partial \text{Vol}_q^\top \mathbf{K}(q, q) \partial \text{Vol}_q} d \text{Vol}_q \right). \end{aligned} \quad (12)$$

Geodesics A geodesic between $q_0, q_1 \in \text{Emb}_1(S_2, \mathbb{R}^3)$ is solution of:

$$\inf \int_0^1 \|v_t\|_K^2 dt \quad (13)$$

under the constraints $\dot{q} = v_t(q(t))$ where $v_t \in V_{q(t)}$. The constraints can be relaxed to $\dot{q} = \pi_{q(t)}(v_t(q(t)))$ where $v_t \in V$. This allows solving Problem (13) by solutions of Hamiltonian equations given by $H(\mu, q) = \frac{1}{2} \langle \pi_q^*(\mu), K(q)\pi_q^*(\mu) \rangle$.

In order to implement these geodesics in the LDDMM framework presented in Section 2.2, we consider that both the control points and the landmark points are part of the system, and write the Hamiltonian

$$H(\mu, q, c) = \frac{1}{2} \mu^T \mathbf{K}(c, c) \mu - \frac{(\mu^T \mathbf{K}(c, q) \partial \text{Vol}_q)^2}{2 \|(d \text{Vol}_q)^\# \|_K^2}.$$

The corresponding Hamiltonian system can readily replace (8) for registration and in the geodesics of the pole ladder constructions. We can thus conduct the same experiment as in the previous section with this metric.

Results We first validate our implementation by measuring the area strain and ejection fraction errors, and complete Table 2 into Table 3. Similarly, we reproduce the boxplot of Figure 15c with the two projection and the previous volume constrained scaling strategy. As in the previous section, there is no relation between the error and the disease group. We show some reconstructions on Figure 19, along the reconstructions obtained by the scaled parallel transport of the previous section. The two metrics, and the two projection methods result in different reconstructions. The projection by gradient flow seems less stable than the two other methods, as the thinning near the valve is exaggerated on the first row, and the apex of the second row is unrealistically spherical.

	Original Values	PT	SPT	grad VPPT	scaling VPPT
AS	-0.24 ± 0.08	0.18	0.13	0.13	0.09
EF	0.42 ± 0.13	0.13	0.04	0.02	0.02

Table 3: Validation of the volume preserving parallel transport (VPPT), with the projection performed either by gradient flow or by scaling.

The values of area strain obtained by this normalization procedure were included in the analysis of Di Folco et al. that leveraged dimensionality reduction methods to compare the different normalization procedures and explore the interactions between shape and deformations. Shape descriptors were computed by point-wise differences between the ED mesh and the atlas after a rigid-body alignment, or alignment by a similarity. By comparing the latent spaces obtained with normalized and unnormalized data, they show that both methods reduce the bias introduced by the alignment of the ED shapes (Di Folco et al. 2021).

Now we come to the original goal of this chapter: performing group-wise analysis of the systolic deformation. We have proposed several ways to normalize each individual time deformation, we now need to summarize the complete systolic trajectory. We propose to use a geodesic - or higher order- regression. This is detailed in the next section dedicated to the downstream analyses.

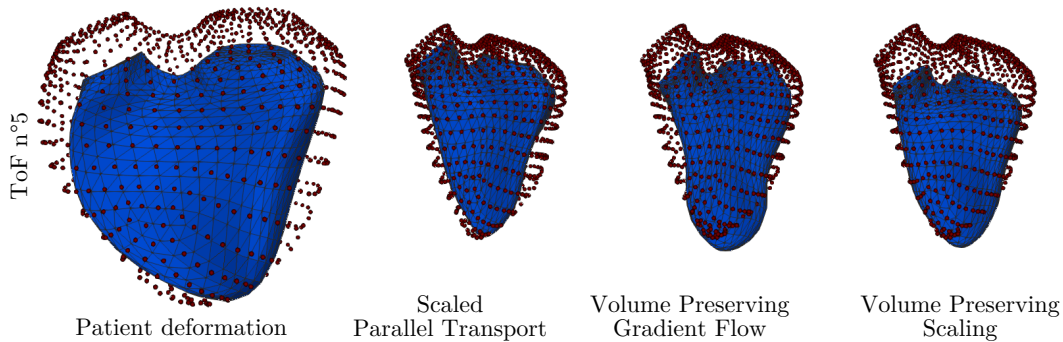


Figure 19: Examples of reconstructions of the ES frame (blue meshes) after parallel transport of the ED-to-ES deformation (point cloud to mesh) along the deformation from ED to atlas. Comparison of the three methods. The gradient flow seems less stable (first two rows) than the projection by scaling. The scaling method results in more vertical movement of the base than the two other methods. This may be due to the centering step required before scaling.

3.5 Analysis of the normalized deformations

In order to compactly represent the full systolic trajectory in the same space and to proceed with linear statistics, we propose to fit a regression model in the shape space as considered in the registration framework of Section 2.2. It estimates a geodesic path between two shapes (the equivalent of a uniform motion). For a trajectory such as the contraction of the cardiac RV, one may expect to find second-order dynamics, making a (first-order) geodesic regression ill-suited. We thus propose to use the second-order model defined in Trouvé and Vialard 2012a to account for the motion of the RV during systole. It is similar to a geodesic regression (Fishbaugh et al. 2017) in the shape space, with an additional acceleration term in the model of the trajectory. We therefore compare it to the simpler geodesic regression.

3.5.1 Geodesic and Spline regression

The model of (Trouvé and Vialard 2012a) introduces a second-order terms u_t can be interpreted as a random external forces smoothly perturbing the trajectory around a mean geodesic which modifies the continuous-time system of equations (8) as follows: $\forall t \in [0, 1]$,

$$\begin{cases} \dot{c}_k^{(t)} &= \sum_j K(c_k^{(t)}, c_j^{(t)}) \mu_j^{(t)} \\ \dot{\mu}_k^{(t)} &= -\sum_j \nabla_1 K(c_k^{(t)}, c_j^{(t)}) \mu_k^{(t)T} \mu_j^{(t)} + u_k^{(t)} \end{cases} \quad (14)$$

Results If we consider a discrete sequence of observation times $t_1 = 0, t_2, \dots, t_d = 1$ and configurations x_{t_1}, \dots, x_{t_d} , one seeks to find the path ϕ_t that minimizes the new cost

$$C_S(c, \mu, u_t) = \frac{1}{\alpha^2 d} \sum_{i=1}^d \|x_{t_i} - \phi_{t_i}(x_{t_0})\|_2^2 + \int_0^1 \|u^{(t)}\|^2 dt + \|v_0^{c, \mu}\|_K^2. \quad (15)$$

In practice, the ODEs (8) and (14) are discretized in n time steps and an integration method such as Euler or Runge-Kutta is used. We define all the patients trajectories

between $t = 0$ and $t = 1$, and use the same discretization for all the patients to ensure that u_0, \dots, u_{N_c} are estimated at corresponding times. Along with $\mu^{(0)}$, these are estimated by gradient descent as in the case of registration. Setting $u_0 = \dots = u_{N_c} = 0$ at all times recovers a geodesic trajectory. We use a kernel bandwidth $\sigma = 15$ in all the experiments, and 60 control points for all the deformations of the atlas. The initial control points are fixed for the entire dataset so that the initial momenta can be compared consistently. They have been optimized to register the atlas on all the transported ES frames.




Figure 20: Fit of the normalized sequence of an ASD patient by geodesic (top) and spline (bottom) regression. Fits are the purple transparent meshes, normalized data are the blue meshes with edges. Upon visual inspection, there is barely any difference between the two fits.

Visually we only notice very slight differences between the geodesic and spline regressions (Figure 20). We do observe that the fit is slightly less precise in the intermediate frames for the geodesic regression, this is confirmed by measuring the overall data attachment term (left term of (15)) that is reported in Table 4, regardless of the normalization method. The spline regression therefore yields a more faithful representation of the normalized cardiac deformations than the geodesic regression, at the cost of a larger set of parameters, encompassing the external time-dependent forces u_t whose interpretation and analysis is difficult. Indeed, the size of the acceleration term u_t is $(d - 1) \times N_c \times 3 = 1620$.

	SPT	grad VPPT	scaling VPPT
Geodesic	1384 ± 952	1659 ± 906	1483 ± 864
Spline	549 ± 413	681 ± 390	554 ± 324

Table 4: Comparison of the mean data attachment term for each method.

3.5.2 Hotelling tests on velocities

The sequences are normalized by one of the three proposed methods: scaled parallel transport (SPT), volume-preserving parallel transport (VPPT) with gradient projection and VPPT with scaling projection, and summarized by a regression with one of the two deformation models: geodesic or spline. In this section, we focus on

the result obtained with the scaled parallel transport and spline regression. This procedure results in a set of descriptive parameters for each individual that lie in a linear space: the co-tangent space at the atlas. For more stable descriptors of the deformation, we compute the initial velocity vectors at the initial control points with $v = \mathbf{K}(c, c)\mu$. We now have all the ingredients to perform statistics on cardiac deformations described by the initial velocities.

We compare the impact of the different diseases by performing pairwise Hotelling tests between each disease group and the control group. One statistic is computed at each control point, that tests if the mean velocity at that point is significantly different. A Bonferroni correction is then applied for multiple testing across the full set of control points, to maintain type I error risk at $\alpha = 0.05$.

The results for the velocities of the spline regression with the scaled PT are displayed Figure 21 for the ASD, ToF and PHT groups and show significant differences between each disease and the control group. The differences are superimposed with

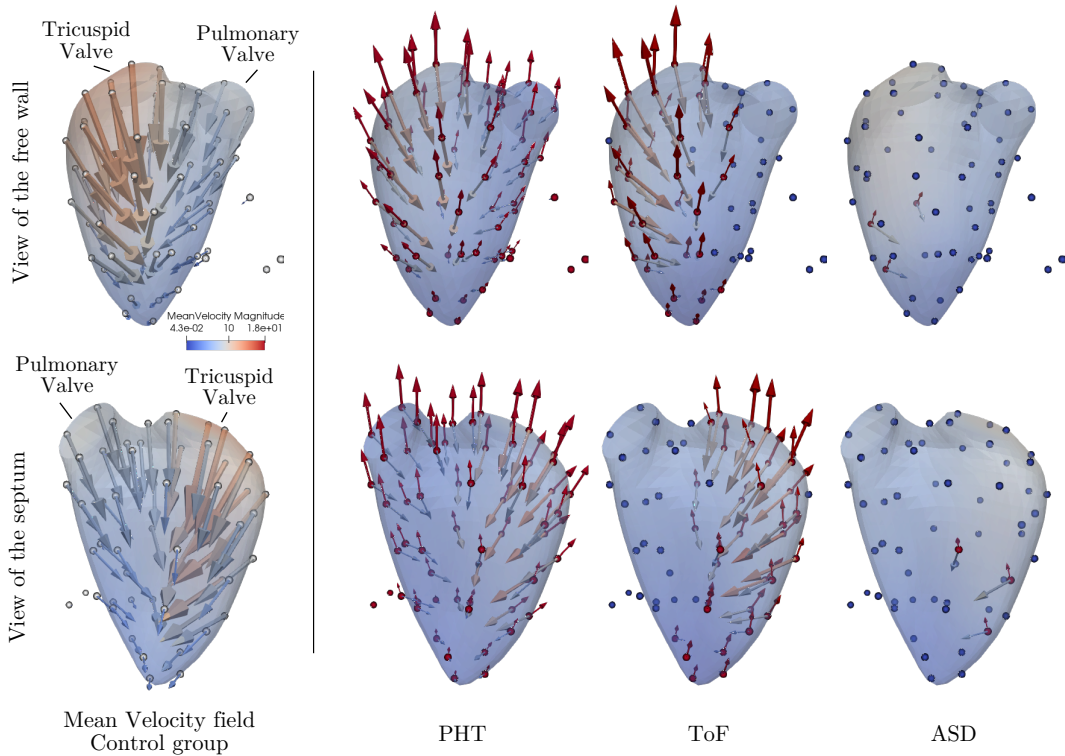


Figure 21: Results of the Hotelling tests between velocities. Where the differences are significant, control points and difference with the control group are in red, while coloured arrows represent the mean velocity field for the disease group. All the arrows have been scaled by a factor 2.3 for visualization purposes. The color map of the meshes reflects the norm of the velocity field at that point, if it is significantly different from the control group. As for the area strain, there is little difference for the ASD group, showing that the volume differences have well been filtered out by the normalization. Moreover the differences for the PHT and ToF group reflect deformations with less amplitude than the control group, as observed on the ejection fraction.

the group mean velocity fields, and show that these differences are mainly of magnitude, with different orientations at a few points along the free wall. The differences observed near the tricuspid valve mainly reflect the difference of magnitude of the deformations, and one should be cautious before drawing further conclusions as the quality of the mesh may vary near this region. However it is interesting to notice that very little differences are observed for the ASD group, which corroborates previous results (Mocerri et al. 2020). Similarly, only small differences are observed on the septum for the PHT group, showing that the shape differences usually observed on the PHT group have been filtered out. This makes the differences observed on the free wall interesting and other markers such as the circumferential strain will be studied to confirm these effects. These experiments also highlight the difference between the PHT and ToF group as significant differences are localized on the inferior part of the free wall and on the inlet (area of the tricuspid valve) for the ToF group while they are distributed across the whole shape for the PHT group.

The same visualization is provided for the mean acceleration term by disease group, across the group-wise mean trajectory (Figure 22). Recall that these can be interpreted as external forces. They seem to increase the movement towards contraction at the beginning of motion while slowing it down towards the end. These terms thus provide more insights into the dynamics of each disease.

Similarly to the analysis of the velocities, the mean external force of the ASD group is not significantly different from that of the control group. This shows that the volume differences reported in Figure 15a were successfully normalized by the scaled parallel transport. For the ToF and PHT group, many points have significantly different values of acceleration, and those differences are shown by red arrows. They are distributed across the whole RV for the PHT group while their locations vary in time for the ToF group, and concentrate around the apex towards the end of the deformation. For the ToF group, we do find differences on the infundibular (area of the pulmonary valve) at times $ED + 1$ and $ED + 2$ as expected from the consequences of the surgery. This is in favor of our method on the analysis of shape deformations, as these differences were not observed with the traditional analysis of strain maps. As in the case of velocities, the differences mainly reflect amplitude differences, especially near ED and ES times. They thus reflect the late contraction of the RV and the distribution of the differences across the ventricle and in time reflect a longer and asynchronous contraction introduced by the disease. Indeed, the deformation of the control group is more uniformly distributed on the shape, and the introduction of heterogeneity across the shape has been previously reported and associated with the disease, and is known to be a factor of arrhythmia with direct consequences on survival.

4 Conclusion

In this chapter, we first summarized recent results using parallel transport for geometric statistics in computational anatomy. We proposed in particular a new implementation of the pole ladder scheme. These results guarantee that parallel transport algorithms are well behaved, enabling their use as a normalization step for longitudinal shape data. This further allows to evaluate the underlying metric and the normalization model beyond the numerical stability.




Figure 22: Results of the Hotelling tests between accelerations. Where the differences are significant, control points and differences with the control group are in red, while coloured arrows represent the mean acceleration for the disease group. As for the velocity fields, there are no differences for the ASD group. For the ToF group, the differences are localized, but these sites vary with time, while for the PHT group, differences are distributed on the entire shape.

On the application side, the results on data of the motion of the cardiac right ventricle exhibited a strong bias due to the volume differences at a reference time point. Two strategies were investigated to correct for this bias: scaling the parallel transport, and using a metric that decomposes volume changes from volume preserving deformations. Both methods were successful at preserving the ejection fraction, and reduced the bias due to volume differences. When scaling the parallel transport, we discovered a significant relationship between the scaling parameter and the initial volume ratio, however its interpretation must be understood further. This will be investigated in future work, together with the possible mathematical relation to the volume preserving metric. The statistical analysis of the normalized deformations was meaningful and coherent with previous knowledge. This framework is readily usable to evaluate the impact of a treatment or a surgery on the deformation. It can also be used to simulate new cardiac sequences for a given shape from a population of cardiac sequences, by performing e.g. a principal component analysis of the momentum vectors of the spline or geodesic regression, and shooting along the principal modes, or by sampling coefficients from a multivariate normal law to mix the principal components.

This study is however not sufficient to decide which method should be preferred. This comes back to the more general question of choosing the metric. As the numerical properties of the parallel transport algorithm are now well grounded, parallel transport can now be used as a proxy to evaluate the pertinence of the metric used. In this implementation of the LDDMM framework, the choice of metric boils down to the choice of kernel, but the framework itself is questioned. As data-driven methods have gained popularity in this field, more papers have proposed strategies to learn the metric from data (Vialard and Risser 2014; Louis et al. 2019; Niethammer, Kwitt, and Vialard 2019). These solutions are attractive to remove the arbitrary aspect of the choice of metric. Yet, beyond the computational complexity, these approaches face the problem of defining ‘natural’ criteria to assess a model. It will be interesting to test the parallel transport normalization method with a metric learned from data.

Acknowledgements This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant G-Statistics agreement No 786854). This work has been supported by the French government, through the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002.

References

- Arsigny, Vincent et al. (2006). “A Log-Euclidean Framework for Statistics on Diffeomorphisms”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*. Ed. by Rasmus Larsen, Mads Nielsen, and Jon Sporring. Vol. 4190. LNCS. Springer Berlin Heidelberg, pp. 924–931. DOI: 10.1007/11866565_113.
- Ashburner, John et al. (1998). “Identifying Global Anatomical Differences: Deformation-based Morphometry”. In: *Human brain mapping* 6.5-6, pp. 348–357.

- Bauer, Martin, Martins Bruveris, and Peter W. Michor (Sept. 2014). “Overview of the Geometries of Shape Spaces and Diffeomorphism Groups”. In: *Journal of Mathematical Imaging and Vision* 50.1-2, pp. 60–97. DOI: 10.1007/s10851-013-0490-z.
- Bône, Alexandre et al. (Sept. 2018). “Deformetrica 4: An Open-Source Software for Statistical Shape Analysis”. In: *Shape in Medical Imaging. ShapeMI 2018*. Ed. by M. Reuter et al. Vol. 11167. LNCS. Springer, pp. 3–13. DOI: 10.1007/978-3-030-04747-4_1.
- Charon, Nicolas et al. (2020). “Fidelity Metrics between Curves and Surfaces: Currents, Varifolds, and Normal Cycles”. In: *Riemannian Geometric Statistics in Medical Image Analysis*. Ed. by Xavier Pennec, Stefan Sommer, and Thomas Fletcher. Academic Press, pp. 441–477. DOI: 10.1016/B978-0-12-814725-2.00021-2.
- Cury, Claire et al. (Sept. 2016). “Spatio-Temporal Shape Analysis of Cross-Sectional Data for Detection of Early Changes in Neurodegenerative Disease”. In: *SeSAMi 2016 - First International Workshop Spectral and Shape Analysis in Medical Imaging*. Vol. 10126. LNCS. Springer, pp. 63–75. DOI: 10.1007/978-3-319-51237-2_6.
- Debavelaere, Vianney et al. (Oct. 2019). “Clustering of Longitudinal Shape Data Sets Using Mixture of Separate or Branching Trajectories”. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*. Vol. Vol 11767. LNCS. Springer, pp. 66–74. DOI: 10.1007/978-3-030-32251-9_8.
- Di Folco, Maxime et al. (2019). “Learning Interactions between Cardiac Shape and Deformation: Application to Pulmonary Hypertension”. In: *10th International Statistical Atlases and Computational Modeling of the Heart (STACOM) Workshop, Held in Conjunction with MICCAI 2019*. Vol. Vol. 12009. LNCS. DOI: 10.1007/978-3-030-39074-7_13.
- Di Folco, Maxime et al. (June 2021). “Investigation of the Impact of Normalization on the Study of Interactions Between Myocardial Shape and Deformation”. In: *Functional Imaging and Modeling of the Heart*. Vol. Vol. 12738. LNCS. Springer, pp. 223–231. DOI: 10.1007/978-3-030-78710-3_22.
- Dryden, I. L. and K. V. Mardia (2016). *Statistical Shape Analysis with Applications in R*. Second edition. Wiley Series in Probability and Statistics. Chichester, UK; Hoboken, NJ: Wiley. ISBN: 978-0-470-69962-1.
- Duchateau, Nicolas et al. (June 2011). “A Spatiotemporal Statistical Atlas of Motion for the Quantification of Abnormal Myocardial Tissue Velocities”. In: *Medical Image Analysis* 15.3, pp. 316–328. DOI: 10.1016/j.media.2010.12.006.
- Durrleman, Stanley, Stéphanie Allasonnière, and Sarang Joshi (Jan. 2013). “Sparse Adaptive Parameterization of Variability in Image Ensembles”. In: *International Journal of Computer Vision* 101.1, pp. 161–183. DOI: 10.1007/s11263-012-0556-1.
- Durrleman, Stanley et al. (May 2013). “Toward a Comprehensive Framework for the Spatiotemporal Statistical Analysis of Longitudinal Shape Data”. In: *International Journal of Computer Vision* 103.1, pp. 22–59. DOI: 10.1007/s11263-012-0592-x.

- Durrleman, Stanley et al. (Nov. 2014). “Morphometry of Anatomical Shape Complexes with Dense Deformations and Sparse Parameters”. In: *NeuroImage* 101, pp. 35–49. DOI: 10.1016/j.neuroimage.2014.06.043.
- Ehlers, Jürgen, Felix A. E. Pirani, and Alfred Schild (1972). “The Geometry of Free Fall and Light Propagation”. In: *General Relativity: Papers in Honour of J. L. Synge*. Ed. by Lochlainn O’Raifeartaigh. Oxford : Clarendon Press, pp. 63–84. ISBN: 978-0-19-851126-7.
- Fishbaugh, James et al. (July 2017). “Geodesic Shape Regression with Multiple Geometries and Sparse Parameters”. In: *Medical Image Analysis* 39, pp. 1–17. DOI: 10.1016/j.media.2017.03.008.
- Gerber, Samuel et al. (Oct. 2010). “Manifold Modeling for Brain Population Analysis”. In: *Medical image analysis* 14.5, pp. 643–653. DOI: 10.1016/j.media.2010.05.008.
- Grenander, U and Michael Miller (1998). “Computational Anatomy: An Emerging Discipline”. In: *Quarterly of Applied Mathematics; LVI(4):617-694*.
- Guigui, Nicolas and Xavier Pennec (June 2021). “Numerical Accuracy of Ladder Schemes for Parallel Transport on Manifolds”. In: *Foundations of Computational Mathematics*. DOI: 10.1007/s10208-021-09515-x.
- Guigui, Nicolas et al. (Apr. 2021). “Cardiac Motion Modeling With Parallel Transport And Shape Splines”. In: *ISBI 2021 - IEEE 18th International Symposium on Biomedical Imaging*, pp. 1394–1397. DOI: 10.1109/ISBI48211.2021.9433887.
- Hadj-Hamou, Mehdi et al. (June 2016). “Longitudinal Analysis of Image Time Series with Diffeomorphic Deformations: A Computational Framework Based on Stationary Velocity Fields”. In: *Frontiers in Neuroscience* 10, p. 236. DOI: 10.3389/fnins.2016.00236.
- Hauberg, Søren, François Lauze, and Kim Steenstrup Pedersen (May 2013). “Unscented Kalman Filtering on Riemannian Manifolds”. In: *Journal of Mathematical Imaging and Vision* 46.1, pp. 103–120. DOI: 10.1007/s10851-012-0372-9.
- Higham, Nicholas J. (Jan. 2005). “The Scaling and Squaring Method for the Matrix Exponential Revisited”. In: *SIAM Journal on Matrix Analysis and Applications* 26.4, pp. 1179–1193. DOI: 10.1137/04061101X.
- Hinkle, Jacob, Thomas Fletcher, and Sarang Joshi (Feb. 2014). “Intrinsic Polynomials for Regression on Riemannian Manifolds”. In: *Journal of Mathematical Imaging and Vision* 50.1-2, pp. 32–52. DOI: 10.1007/s10851-013-0489-5.
- Kheyfets, Arkady, Warner A. Miller, and Gregory A. Newton (Dec. 2000). “Schild’s Ladder Parallel Transport Procedure for an Arbitrary Connection”. In: *International Journal of Theoretical Physics* 39.12, pp. 2891–2898. DOI: 10.1023/A:1026473418439.
- Kleijn, Sebastiaan A. et al. (Mar. 2011). “Three-Dimensional Speckle Tracking Echocardiography for Automatic Assessment of Global and Regional Left Ventricular Function Based on Area Strain”. In: *Journal of the American Society of Echocardiography* 24.3, pp. 314–321. DOI: 10.1016/j.echo.2011.01.014.
- Lorenzi, Marco, Nicholas Ayache, and Xavier Pennec (July 2011). “Schild’s Ladder for the Parallel Transport of Deformations in Time Series of Images”. In: *IPMI - 22nd International Conference on Information Processing in Medical Images-2011*. Vol. 6801. Springer, p. 463. DOI: 10.1007/978-3-642-22092-0_38.

- Lorenzi, Marco and Xavier Pennec (Nov. 2013). “Geodesics, Parallel Transport & One-parameter Subgroups for Diffeomorphic Image Registration”. In: *International Journal of Computer Vision* 105.2, pp. 111–127. DOI: 10.1007/s11263-012-0598-4.
- (Sept. 2014). “Efficient Parallel Transport of Deformations in Time Series of Images: From Schild to Pole Ladder”. In: *Journal of Mathematical Imaging and Vision* 50.1, pp. 5–17. DOI: 10.1007/s10851-013-0470-3.
- Lorenzi, Marco et al. (Sept. 2011). “Mapping the Effects of a β Levels on the Longitudinal Changes in Healthy Aging: Hierarchical Modeling Based on Stationary Velocity Fields”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2011*. Springer, Berlin, Heidelberg, pp. 663–670. DOI: 10.1007/978-3-642-23629-7_81.
- Lorenzi, Marco et al. (Jan. 2015). “Disentangling Normal Aging from Alzheimer’s Disease in Structural Magnetic Resonance Images”. In: *Neurobiology of Aging* 36, S42. DOI: 10.1016/j.neurobiolaging.2014.07.046.
- Louis, Maxime et al. (2017). “Parallel Transport in Shape Analysis: A Scalable Numerical Scheme”. In: *Geometric Science of Information*. Ed. by Frank Nielsen and Frédéric Barbaresco. Vol. 10589. Springer International Publishing, pp. 29–37. DOI: 10.1007/978-3-319-68445-1_4.
- Louis, Maxime et al. (2018). “A Fanning Scheme for the Parallel Transport Along Geodesics on Riemannian Manifolds”. In: *SIAM Journal on Numerical Analysis* 56.4, pp. 2563–2584. DOI: 10.1137/17M1130617.
- Louis, Maxime et al. (June 2019). “Riemannian Geometry Learning for Disease Progression Modelling”. In: *Information Processing in Medical Imaging*. Springer, Cham, pp. 542–553. DOI: 10.1007/978-3-030-20351-1_42.
- Mansi, Tommaso et al. (2009). “A Statistical Model of Right Ventricle in Tetralogy of Fallot for Prediction of Remodelling and Therapy Planning”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2009*. Ed. by Guang-Zhong Yang et al. Vol. 5761. Series Title: Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 214–221. DOI: 10.1007/978-3-642-04268-3_27.
- McLeod, Kristin et al. (2013). “Statistical Shape Analysis of Surfaces in Medical Images Applied to the Tetralogy of Fallot Heart”. In: *Modeling in Computational Biology and Biomedicine: A Multidisciplinary Endeavor*. Ed. by Frédéric Cazals and Pierre Kornprobst. Berlin, Heidelberg: Springer, pp. 165–191. DOI: 10.1007/978-3-642-31208-3_5.
- Micheli, Mario and Joan A. Glaunès (2014). “Matrix-Valued Kernels for Shape Deformation Analysis”. In: *Geometry, Imaging and Computing* 1.1, pp. 57–139. DOI: 10.4310/GIC.2014.v1.n1.a2.
- Miller, Michael, Alain Trounev, and Laurent Younes (Dec. 2015). “Hamiltonian Systems and Optimal Control in Computational Anatomy: 100 Years Since D’Arcy Thompson”. In: *Annual Review of Biomedical Engineering* 17.1, pp. 447–509. DOI: 10.1146/annurev-bioeng-071114-040601.
- Misner, Charles W., Kip S. Thorne, and John Archibald Wheeler (1973). *Gravitation*. Ed. by Freeman. Princeton University Press. ISBN: 978-0-691-17779-3.
- Mocerì, Pamela et al. (Apr. 2018). “Three-Dimensional Right-Ventricular Regional Deformation and Survival in Pulmonary Hypertension”. In: *European Heart*

- Journal - Cardiovascular Imaging* 19.4, pp. 450–458. DOI: 10.1093/ehjci/jex163.
- Moceri, Pamela et al. (2020). “3D Right Ventricular Shape and Strain in Congenital Heart Disease Patients with Right Ventricular Chronic Volume Loading”. In: *European Heart Journal - Cardiovascular Imaging* 22.10. DOI: 10.1093/ehjci/jeaa189.
- Niethammer, Marc, Roland Kwitt, and François-Xavier Vialard (2019). “Metric Learning for Image Registration”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8463–8472.
- Niethammer, Marc and François-Xavier Vialard (2013). “Riemannian Metrics for Statistics on Shapes : Parallel Transport and Scale Invariance”. In: *Proceedings of Miccai Workshop, MFCA*.
- Pennec, Xavier (June 2019). *Curvature Effects on the Empirical Mean in Riemannian and Affine Manifolds: A Non-Asymptotic High Concentration Expansion in the Small-Sample Regime*. Tech. rep. arXiv: 1906.07418.
- Pennec, Xavier and Vincent Arsigny (May 2012). “Exponential Barycenters of the Canonical Cartan Connection and Invariant Means on Lie Groups”. In: *Matrix Information Geometry*. Springer, pp. 123–168. DOI: 10.1007/978-3-642-30232-9_7.
- Pennec, Xavier and Marco Lorenzi (Jan. 2020). “Beyond Riemannian Geometry: The Affine Connection Setting for Transformation Groups”. In: *Riemannian Geometric Statistics in Medical Image Analysis*. Ed. by Xavier Pennec, Stefan Sommer, and Thomas Fletcher. Academic Press, pp. 169–229. DOI: 10.1016/B978-0-12-814725-2.00012-1.
- Pennec, Xavier, Stefan Sommer, and Thomas Fletcher, eds. (2020). *Riemannian Geometric Statistics in Medical Image Analysis*. The Elsevier and MICCAI Society Book Series. Elsevier. ISBN: 978-0-12-814725-2. DOI: 10.1016/C2017-0-01561-6.
- Peyrat, Jean-Marc et al. (July 2010). “Registration of 4D Cardiac CT Sequences Under Trajectory Constraints With Multichannel Diffeomorphic Demons”. In: *IEEE Transactions on Medical Imaging* 29.7, pp. 1351–1368. DOI: 10.1109/TMI.2009.2038908.
- Qiu, Anqi et al. (Mar. 2008). “Parallel Transport in Diffeomorphisms Distinguishes the Time-Dependent Pattern of Hippocampal Surface Deformation Due to Healthy Aging and the Dementia of the Alzheimer’s Type”. In: *NeuroImage* 40.1, pp. 68–76. DOI: 10.1016/j.neuroimage.2007.11.041.
- Qiu, Anqi et al. (Mar. 2009). “Time Sequence Diffeomorphic Metric Mapping and Parallel Transport Track Time-Dependent Shape Changes”. In: *NeuroImage. Mathematics in Brain Imaging* 45.1, Supplement 1, S51–S60. DOI: 10.1016/j.neuroimage.2008.10.039.
- Rao, A. et al. (Sept. 2004). “Spatial Transformation of Motion and Deformation Fields Using Nonrigid Registration”. In: *IEEE Transactions on Medical Imaging* 23.9, pp. 1065–1076. DOI: 10.1109/TMI.2004.828681.
- Sanz, Javier et al. (Apr. 2019). “Anatomy, Function, and Dysfunction of the Right Ventricle: JACC State-of-the-Art Review”. In: *Journal of the American College of Cardiology* 73.12, pp. 1463–1482. DOI: 10.1016/j.jacc.2018.12.076.

- Schiratti, Jean-Baptiste et al. (2015). “Learning Spatiotemporal Trajectories from Manifold-Valued Longitudinal Data”. In: *Advances in Neural Information Processing Systems* 28.
- Singh, Nikhil, François-Xavier Vialard, and Marc Niethammer (Oct. 2015). “Splines for diffeomorphisms”. In: *Medical Image Analysis* 25.1, pp. 56–71. DOI: 10.1016/j.media.2015.04.012.
- Sivera, Raphaël et al. (Oct. 2020). “Voxel-Based Assessments of Treatment Effects on Longitudinal Brain Changes in the Multidomain Alzheimer Preventive Trial Cohort”. In: *Neurobiology of Aging* 94, p. 50. DOI: 10.1016/j.neurobiolaging.2019.11.020.
- Thompson, D’Arcy Wentworth (1917). *On Growth and Form*. Cambridge: Cambridge University Press. DOI: 10.1017/CB09781107325852.
- Trouvé, Alain and François-Xavier Vialard (2012a). “Shape Splines and Stochastic Shape Evolutions: A Second Order Point of View”. In: *Quarterly of Applied Mathematics* 70.2, pp. 219–251.
- (2012b). “Shape splines and stochastic shape evolutions: a second order point of view”. In: *Quarterly of Applied Mathematics* 70.2. Publisher: Brown University, pp. 219–251. URL: <https://www.jstor.org/stable/43639026>.
- Vialard, François-Xavier and Laurent Risser (Sept. 2014). “Spatially-Varying Metric Learning for Diffeomorphic Image Registration: A Variational Framework”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2014*. Vol. 8673. LNCS. Cham: Springer, pp. 227–234. DOI: 10.1007/978-3-319-10404-1_29.
- Younes, Laurent (Feb. 2007). “Jacobi Fields in Groups of Diffeomorphisms and Applications”. In: *Quarterly of Applied Mathematics* 65.1, pp. 113–134. DOI: 10.1090/S0033-569X-07-01027-5.
- (2019). *Shapes and Diffeomorphisms*. Vol. 171. Applied Mathematical Sciences. Berlin, Heidelberg: Springer. ISBN: 978-3-662-58495-8. DOI: 10.1007/978-3-662-58496-5.
- Young, Alistair A. and Alejandro F. Frangi (2009). “Computational Cardiac Atlases: From Patient to Population and Back”. In: *Experimental Physiology* 94.5, pp. 578–596. DOI: 10.1113/expphysiol.2008.044081.