

Hybrid Riemannian metrics for diffeomorphic shape registration

LAURENT YOUNES^{*,†}

We consider the results of combining two approaches developed for the design of Riemannian metrics on curves and surfaces, namely parametrization-invariant metrics of the Sobolev type on spaces of immersions, and metrics derived through Riemannian submersions from right-invariant Sobolev metrics on groups of diffeomorphisms (the latter leading to the “large deformation diffeomorphic metric mapping” framework). We show that this quite simple approach inherits the advantages of both methods, both on the theoretical and experimental levels, and provide additional flexibility and modeling power, especially when dealing with complex configurations of shapes. Experimental results illustrating the method are provided for curve and surface registration.

AMS 2000 SUBJECT CLASSIFICATIONS: 58D05, 49N90, 49Q10, 68E10.

KEYWORDS AND PHRASES: Shape analysis, groups of diffeomorphisms, sub-Riemannian metrics, optimal control, computational anatomy.

1. Introduction

1.1. Shape registration: basic principles

We consider a “shape space” (denoted \mathcal{M}) in which objects are subject to free-form deformations, so that a group action $(\varphi, q) \in \text{Diff} \times \mathcal{M} \mapsto \varphi \cdot q \in \mathcal{M}$ is defined on \mathcal{M} (where Diff refers to the diffeomorphism group). This concept usually comes with additional conditions on the structure of the space. Here, following [1], we will assume that \mathcal{M} is an open subset of a Banach space \mathcal{Q} . One also often considers \mathcal{M} quotiented by other group actions (such as Euclidean transformations, or reparametrization). Even though we will not formally consider such quotient spaces, such invariance will often be a direct consequence of the models we will discuss.

One can interpret registration methods within the following framework. Given a “template” $q_0 \in \mathcal{M}$, a registration method can be seen as a trans-

^{*}This paper is dedicated to David Mumford on his 80th birthday.

[†]Partially supported by NIH U19AG033655, R01HL130292 and R01DC016784.

formation that takes a shape q as input and returns as output a diffeomorphism φ such that $\varphi \cdot q_0 = (\text{or } \simeq) q$, therefore providing a mapping, $\Psi_{q_0} : \mathcal{M} \rightarrow \text{Diff}$, representing each shape by a diffeomorphism. One of the main advantages of such constructions is that it is much easier to define characteristic features of diffeomorphisms than of shapes, considered, for example, as subsets of \mathbb{R}^d . Voxel-based, or surface-based morphometric methods have exploited this by introducing deformation markers, often deduced from the Jacobian of the estimated diffeomorphism [4, 32].

When working with shape spaces of landmarks, images, curves or surfaces, there exist, for each given shape q , either zero or an infinity of transformations such that $\varphi \cdot q_0 = q$. They can all be deduced from each other via composition on the right by diffeomorphisms that leave q_0 invariant, i.e., elements of the stabilizer $\mathcal{S}_{q_0} = \{\psi \in \text{Diff} : \psi \cdot q_0 = q_0\}$. “Good” registration algorithms generally pick one such transformation that maximizes a regularity or optimality criterion that the algorithm implements. Understanding the structure of the space $\Psi_{q_0}(\Omega)$ has many advantages, because it may lead to (locally) one-to-one shape representations. Among others methods [16, 31, 17, 35, 34, 5, 19, 36, 25, 21, 6], the large deformation diffeomorphic metric mapping framework (LDDMM) includes a family of registration algorithms [22, 23, 7, 12, 13, 18, 14, 33], adapted to various shape modalities, that rely on a sub-Riemannian setup of the diffeomorphism group and of the shape space. This setup is a special case of the model used in this paper, that we now describe.

We will assume a sub-Riemannian structure on Diff and consider conditions under which it can be transferred into a sub-Riemannian structure on \mathcal{M} through the action. This framework will include the LDDMM construction as a special case, and the other metrics that will be used in this paper. The following notation and assumptions will be used throughout this paper.

We will only consider diffeomorphisms that tend to the identity at infinity, i.e., $\varphi = \text{id} + u$ such that u and du tend to 0 at infinity. Diff will denote the space of such diffeomorphisms, and $u \mapsto \text{id} + u$ provides a trivial chart of Diff as a Banach manifold, when defined over the space $C_0^1(\mathbb{R}^d, \mathbb{R}^d)$ of u 's that tend to 0 at infinity (with the supremum norm: $\|u\|_{1,\infty} = \|u\|_\infty + \|du\|_\infty$). We will assume that the action $(\varphi, q) \mapsto \varphi \cdot q$ is C^1 from $\text{Diff} \times \mathcal{M} \rightarrow \mathcal{M}$. We let $\pi_q(\varphi) = \varphi \cdot q$ and $\xi_q = d\pi_q(\text{id})$, so that the infinitesimal action is given by $v \cdot q = \xi_q v$.

Let V be a Hilbert space continuously embedded in $C_0^1(\mathbb{R}^d, \mathbb{R}^d)$. Consider the sub-Riemannian structure on Diff associated with the distribution $V_\varphi = \{v \circ \varphi : v \in V\}$. We assume that V_φ is equipped with a Hilbert structure, with norm $\|\cdot\|_\varphi$ such that, for all $v \in V$, $\|v \circ \varphi\|_\varphi \geq c\|v\|_V$ for some $c > 0$. We will also denote $\|v\|_{V,\varphi} = \|v \circ \varphi\|_\varphi$ for $v \in V$. A path $(\varphi(t), t \in [0, 1])$ in

$Diff$ is admissible if $\partial_t \varphi \in V_\varphi$ for (almost all) t and

$$\int_0^1 \|\partial_t \varphi\|_{\varphi(t)}^2 dt < \infty.$$

(Paths tangent to $V_{\varphi(t)}$ at all times are often called horizontal in the sub-Riemannian literature. We will not use this term here because of the horizontal spaces that we define just below.)

Let $Diff_0$ be the subgroup of $Diff$ containing all the elements that are reachable from the identity with an admissible path. Fix $q_0 \in \mathcal{M}$ and let $\mathcal{M}_0 = \{\pi_{q_0}(\varphi) : \varphi \in Diff_0\}$. For $\varphi \in Diff_0$ and $q = \pi_{q_0}(\varphi)$, consider the space orthogonal to $Null(\xi_q)$ for $\langle \cdot, \cdot \rangle_{V,\varphi}$, denoted H_φ . Assume that the Hilbert space isometry

$$(1) \quad (H_\varphi, \|\cdot\|_{V,\varphi}) \sim (H_\psi, \|\cdot\|_{V,\psi})$$

holds whenever $\pi_{q_0}(\varphi) = \pi_{q_0}(\psi)$. Then, for $\pi_{q_0}(\varphi) = q$, one can define the space $\mathcal{H}_q = \xi_q H_\varphi = \{\xi_q v : v \in H_\varphi\}$ with

$$\|\xi_q v\|_q = \|v\|_{V,\varphi}$$

and this definition does not depend on which $\varphi \in \pi_{q_0}^{-1}(q)$ is chosen. The distribution \mathcal{H}_q then provides a sub-Riemannian structure on \mathcal{M}_0 .

The space $\mathcal{H}_\varphi = \{v \circ \varphi : v \in H_\varphi\}$ is the horizontal space at φ for the mapping π_{q_0} and the statement that these spaces are isometric adapts the conditions for π_{q_0} to be a Riemannian submersion to this sub-Riemannian setting. In the LDDMM framework, (1) is ensured by defining $\|v \circ \varphi\|_\varphi = \|v\|_V$ for all φ and $v \in V$, so that the original metric is right-invariant. We will below consider computationally feasible situations in which the latter condition is relaxed with (1) still holding.

From a practical viewpoint, LDDMM can be expressed as an optimal control problem. One can indeed describe the search for a geodesic between the template q_0 and a shape $q_1 \in \mathcal{M}$ as the minimization, over all time-dependent vector fields $v : [0, 1] \rightarrow V$, of

$$(2) \quad \frac{1}{2} \int_0^1 \|v(t)\|_V^2 dt,$$

subject to $\partial_t q(t) = v(t) \cdot q(t)$, $q(0) = q_0$, and $q(1) = q_1$. LDDMM uses a relaxation of the last constraint, minimizing

$$(3) \quad \frac{1}{2} \int_0^1 \|v(t)\|_V^2 dt + D(q(1), q_1)$$

subject to $\partial_t q(t) = v(t) \cdot q(t)$ and $q(0) = q_0$, where $D(q, \tilde{q})$ is some properly defined discrepancy measure between q and \tilde{q} . Invariance is often ensured by considering functions such that $D(q, \tilde{q}) = 0$ if q and \tilde{q} are related by a transformation for which the invariance is searched.

Because of the embedding assumption, the norm on V is associated with a reproducing kernel, which is a matrix-valued function:

$$\begin{aligned} K : \mathbb{R}^d \times \mathbb{R}^d &\rightarrow \mathcal{M}_d(\mathbb{R}^d) \\ (x, y) &\mapsto K(x, y) \end{aligned}$$

such that, for all $x, a \in \mathbb{R}^d$, $K(\cdot, x)a : y \mapsto K(y, x)a$ belongs to V , and for all $v \in V$,

$$a^T v(x) = \langle v, K(\cdot, x)a \rangle_V.$$

This kernel and the norm on V are generally chosen to be translation invariant, taking the form

$$(4) \quad K(x, y) = \Gamma\left(\frac{x - y}{a}\right)$$

where Γ is a positive definite (matrix-valued) function. The extra parameter, a , can be interpreted as a scale parameter, that can be tuned to modulate the locality of the deformations. It essentially modulates the long range effect of the motion of a single particle in space. For example, the kernel associated with

$$\|v\|_V^2 = \int_{\mathbb{R}^d} |(\text{Id} - a^2 \Delta)^{m/2} v|^2 dx,$$

where $c := m - (d + 1)/2$ is a positive integer, is given by (4) with

$$(5) \quad \Gamma(x) = P_c(|x|)e^{-|x|},$$

where P_c is a *reverse Bessel polynomial* of degree c (see [26]), normalized so that $P(0) = 1$. The associated kernel K decays to 0 at infinity, at a speed which is modulated by the scale constant a . The shape of the function Γ for $a = 1$ and $c = 0, 1, \dots, 4$ is provided in Figure 1 (Γ has c continuous derivatives at $x = 0$ and is C^∞ everywhere else). We used $c = 3$ in our experiments. For this kernel, the half-range (value of $|x|$ for which $\Gamma(x) = 1/2$) is given by $2.85a$.

The rather simple formulation (ignoring numerical issues) provides a horizontal geodesic in Diff_0 , given by the flow associated with an optimal v , i.e., the solution of $\partial_t \varphi(t) = v(t) \circ \varphi(t)$ with initial condition $\varphi(0) = \text{id}$.




Figure 1: Shapes of the kernels given by (5) for various values of c .

2. Hybrid LDDMM

Starting with a right-invariant metric on $Diff_0$ ensures that all tangent spaces are isometric to the tangent space at the identity, through the right-translation map $v \mapsto v \circ \varphi$. This property is much stronger than what is needed to ensure (1), which only requires that horizontal spaces within the same fibers to be isometric. Obviously, right-invariance brings additional properties to the Riemannian structure on $Diff_0$, making it, in particular, independent of the choice of the template, q_0 , and ensuring that its geodesics satisfy strong conservation laws [20, 41, 30, 40]. On the other hand, it prevents the metric from taking into account shape-dependent properties, related, for example to the geometry of the considered curves or surfaces.

Allowing for less restrictive invariance will allow us to characterize a much larger variety of features compared to those associated with plain LDDMM. This is done in the following examples over spaces of curves and surfaces.

2.1. Curves

2.1.1. Hybrid norms. Consider the situation in which the objects of interest are curves, in two or three dimensions. In this setting, we can take advantage of the collection of metrics that have been introduced for shape spaces of embedded curves $q : M \rightarrow \mathbb{R}^2$, where M is either the unit interval or the unit disk. We let $\mathcal{Q} = C^r(M, \mathbb{R}^d)$ (normed by the supremum norm

over all derivatives of order r or less), for some $r > 1$, and \mathcal{M} be the set of C^r embeddings. The action being $\varphi \cdot q = \varphi \circ q$, we have $\xi_q v = v \circ q$, and we will assume that V is continuously embedded in $C_0^p(\mathbb{R}^d, \mathbb{R}^d)$ with $p > r$, which ensures that $(v, q) \mapsto \xi_q v \in \mathcal{Q}$ is C^1 .

Several Riemannian metrics on such shape spaces have been introduced and studied in the literature (see [39, 24, 37, 28, 29, 38, 9]), most of the times defined via reparametrization-independent positive self-adjoint differential operators, L_q , leading to norms taking the form $\|h\|_q^2 = \int_M h(x)^T (L_q h)(x) |q'(x)| dx$. For example, one can use

$$(6) \quad \|h\|_q^2 = \int_M (\alpha |h(x)|^2 |q'(x)| + |h'(x)|^2 |q'(x)|^{-1}) dx,$$

for which $L_q h = \alpha h - |q'|^{-1} \partial_x (h' |q'|^{-1}) = \alpha h - \partial_s^2 h$. Here and in the following, we use either $\partial_x f$ or f' to denote the derivative of a function with respect to x (assuming no ambiguity on the variable in the second case) and ∂_s , the arc-length derivative, denotes the operator $|q'|^{-1} \partial_x$. Higher-order derivatives have also been studied, together with the introduction of weights relying on geometric properties like length or curvature. We will always assume that there exists a function C defined on \mathcal{M} such that $\|h\|_q \leq C(q) \|h\|_{\mathcal{Q}}$ for all $h \in \mathcal{Q}$ and $q \in \mathcal{M}$ and that $q \mapsto \|h\|_q^2$ is C^1 from \mathcal{M} to \mathbb{R} .

Choosing one of these metrics, $h \mapsto \|h\|_q$, applied to vector fields along q , we introduce the norm, applied to vector fields v defined over \mathbb{R}^d :

$$(7) \quad \|v\|_q^2 = \lambda \|v\|_V^2 + \|v \circ q\|_q^2.$$

It is important to notice that $\|\cdot\|_q$ can be a semi-norm in this expression (i.e., one can have $\|h\|_q = 0$ and $h \neq 0$), while still ensuring that $\|\cdot\|_q$ is a norm. For example, one can take $\alpha = 0$ in (6), which has the nice property of making this semi-norm blind to translations (for which h is constant). In our experiments, we use a version of this norm which is, in addition, blind to rotations, given by (as a function of the arc length)

$$(8) \quad \|h\|_q^2 = \int_0^{\ell(q)} |\partial_s h|^2 ds - \frac{1}{\ell(q)} \left(\int_0^{\ell(q)} \partial_s h^T N(s) ds \right)^2$$

where $\ell(q) = \int_M |\dot{q}| dx$ is the length of q and $N(s)$ is the unit normal to q at $q(s)$. Another interesting norm [39, 38] is

$$\|h\|_q^2 = \frac{1}{\ell(q)} \int_0^{\ell(q)} |\partial_s h|^2 ds$$

and its corresponding rotation/scale-invariant version

$$(9) \quad \begin{aligned} \llbracket h \rrbracket_q^2 &= \frac{1}{\ell(q)} \int_0^{\ell(q)} |\partial_s h|^2 ds \\ &\quad - \left(\frac{1}{\ell(q)} \int_0^{\ell(q)} \partial_s h^T T(s) ds \right)^2 - \left(\frac{1}{\ell(q)} \int_0^{\ell(q)} \partial_s h^T N(s) ds \right)^2 \end{aligned}$$

where $T(s)$ is the unit tangent.

Given such a norm, we can consider what we refer to as an *hybrid LD-DMM* problem minimizing

$$(10) \quad \int_0^1 \|v(t)\|_{q(t)}^2 dt + D(q(1), q_1)$$

subject to $q(0) = q_0$ and $\partial_t q(t) = v(t) \circ q(t)$.

There are two ways to interpret this approach. The first one, following our presentation, is that it modifies LDDMM by taking into account geometric properties of the curve. Alternatively, one can interpret this norm as a modification of one of the norms used on spaces of immersed curves, who generally do not prevent self intersections along geodesic paths. From this viewpoint, the first term ($\|v\|_V^2$) is a global control ensuring the existence of a diffeomorphism of \mathbb{R}^d transforming the curve, therefore guaranteeing that the curve remains embedded along any finite-energy trajectory.

One of the most interesting applications of this formulation is that it is easy to generalize it to the comparison of multiple curves, say q_1, \dots, q_n , using

$$(11) \quad \|v\|_q^2 = \lambda \|v\|_V^2 + \llbracket v \circ q_1 \rrbracket_{q_1}^2 + \dots + \llbracket v \circ q_n \rrbracket_{q_n}^2.$$

Assuming that the curves do not overlap to start with, they will remain apart along any trajectory. However, if one chooses a small scale for the kernel of V , these curves will have almost no interaction unless they come close to each other. The terms $\llbracket \cdot \rrbracket_{q_i}$ control the shape variations of each curve separately, while $\|v\|_V$ ensures global consistency via a diffeomorphic transformation. With such a model, it becomes possible, by playing with the permissiveness of the V -norm relative to the curve metrics, to transform sets of curves while ensuring that each curve evolves in an almost rigid way while their relative position in space may vary greatly. This is illustrated by several examples in section 2.3.

Note that multiple shape comparison using the LDDMM approach was also the subject of [3]. The approach in that work differs from what we are proposing here, because in [3], each curve was attributed its own vector field, say v_i , separately controlled by an RKHS norm, and global consistency was ensured by an additional vector field (similar to the v that we are using here) and by equality constraints for the curve displacement, ensuring that $v \circ q = v_i \circ q$. The approach in the present paper is significantly easier to implement, avoiding, in particular, the need for constrained optimization methods.

2.2. Maximum principle and optimization algorithm

Consider the Hamiltonian defined for $(p, q) \in \mathcal{Q}^* \times \mathcal{Q}$ and $v \in V$, by

$$H_v(p, q) = (p \mid v \circ q) - \frac{1}{2} \|v\|_q^2$$

where $(\cdot \mid \cdot)$ denotes the pairing between elements of a vector space and of its dual space, i.e., the evaluation of the first variable on the second one. If one adds the assumption that $q \mapsto D(q, q_1)$ is differentiable from \mathcal{M} to \mathbb{R} , one can prove that the Pontryagin's maximum principle (PMP) applies. This principle states that solutions of (10) are such that there exists $p : [0, 1] \rightarrow \mathcal{Q}^*$ satisfying

$$(12) \quad \begin{cases} \partial_t q = \partial_p H_{v(t)}(p, q) \\ \partial_t p = -\partial_q H_{v(t)}(p, q) \\ v(t) = \operatorname{argmax}_w H_w(p(t), q(t)) \end{cases}$$

with the boundary conditions $q(0) = q_0$ and $p(1) = -\partial_q D(q, q_1)$. The validity of the principle can be derived from the differentiability of the equation $\partial_t q = v \circ q$ with respect to the control (v) and several applications of the chain rule. (We skip the proof.)

When $\|v\|_q$ takes the form in (7), the Hamiltonian, considered as a function of v takes the form $\lambda \|v\|_V^2 + F_{p,q}(\xi_q v)$ for a C^1 function $F_{p,q} : \mathcal{Q} \rightarrow \mathbb{R}$. As a consequence, the optimal v in the third equation of (12) is such that $\lambda \langle v, h \rangle_V + (dF_{p,q} \mid \xi_q h) = 0$ for all $h \in V$, and writing

$$(dF_{p,q} \mid \xi_q h) = (\xi_q^* dF_{p,q} \mid h) = \langle K \xi_q^* dF_{p,q}, h \rangle_V$$

we see that v should take the form $v = K \xi_q^* \alpha$ for some $\alpha \in \mathcal{Q}^*$. One can therefore apply the same reduction as the one which is typically used with standard LDDMM [2], using α as a new control and minimizing

$$(13) \quad \frac{1}{2} \int_0^1 \|\alpha(t)\|_{q(t)}^2 dt + D(q(1), q_1)$$

subject to $q(0) = q_0$ and $\partial_t q(t) = \mathcal{K}_{q(t)}\alpha(t)$, with $\mathcal{K}_q = \xi_q K \xi_q^*$ and

$$\|\alpha\|_q^2 = \lambda(\alpha \mid \mathcal{K}_q \alpha) + \|\mathcal{K}_q \alpha\|_q^2.$$

The PMP can then be rewritten starting from the Hamiltonian

$$H_\alpha(p, q) = (p \mid \mathcal{K}_q \alpha) - \frac{1}{2} \|\alpha\|_q^2.$$

In the case we are considering in this paragraph, for which $\xi_q v = v \circ q$, \mathcal{K}_q is given by

$$(\mathcal{K}_q \alpha)(x) = (\alpha \mid K(q(x), q(\cdot))).$$

In our numerical implementations, in which q is represented as polygon with vertexes (x_1, \dots, x_N) , $\|\alpha\|_q^2$ is approximated using finite differences, so that the approximation is a function of $(h(x_i), i = 1, \dots, N)$ and of (x_1, \dots, x_N) . As a result, the optimal control problem is reduced to a problem where state and controls are in $(\mathbb{R}^d)^N$ (and the metric is Riemannian). The numerical results that we provide are based on this approximation (and a time approximation using a standard Euler scheme). We also recall that the computation of the gradient of the objective function (considered as a function of the control) can be based on the PMP, using the adjoint algorithm that first computes $q(\cdot)$ by solving the first equation in (12) starting with q_0 and using the control at which the gradient is computed; then sets $p(1) = -\partial_q D(q, q_1)$ before solving the second equation backward in time to obtain $p(\cdot)$; and finally computes the differential of the objective function, which is given by $\partial_v H_v(p, q)$ (or $\partial_\alpha H_\alpha(p, q)$). The gradient itself is defined by $\mathcal{K}_q^{-1} \partial_\alpha H_\alpha(p, q)$, which requires no operator inversion, because

$$\mathcal{K}_q^{-1} \partial_\alpha H_\alpha(p, q) = p - \alpha - \mathcal{G}_q \mathcal{K}_q \alpha$$

where $\mathcal{G}_q : \mathcal{Q} \rightarrow \mathcal{Q}^*$ is the operator defined by $\langle h_1, h_2 \rangle_q = (\mathcal{G}_q h_1 \mid h_2)$.

2.3. Experiments

Cost function. The end-point cost we used for our experiments is a version of the varifold metric introduced in [15]. More precisely, we let

$$(14) \quad D(q, q_1) = \|q\|_X^2 - 2\langle q, q_1 \rangle_X + \|q_1\|_X^2$$




Figure 2: Smoothed Cardioids: Estimated deformations. (a) Template (red) and target (blue); (b) Standard LDDMM; (c) Hybrid LDDMM.

where

(15)

$$\langle q, q_1 \rangle_\chi^2 = \int_M \int_M \chi(q(u), q_1(u_1)) (1 + c(\nu(u)^T \nu_1(u_1))^2) |q'(u)| |q'_1(u_1)| du du_1,$$

where ν and ν_1 denote the unit normals to q and q_1 , χ is a Gaussian kernel

$$\chi(x, y) = \exp(-|x - y|^2 / 2\tau^2),$$

τ and c being fixed parameters (we used $\tau = 2$ and $c = 1$ in our experiments). Because this cost function is bi-invariant by reparametrization ($D(q, q \circ \beta) = 0$ if β is a diffeomorphism of M) and $\langle \cdot, \cdot \rangle_q$ is invariant too ($\langle h \circ \beta \rangle_{q \circ \beta} = \langle h \rangle_q$), the problem is reparametrization-invariant (replacing q_1 by $q_1 \circ \beta$ does not change the solution).

2.3.1. Smoothed cardioids. We first illustrate the impact of the additional energy term with a simple example in which two smoothed curves are registered (see right panel in Figure 2). We used standard LDDMM with a kernel size $a = .2$ (the size of the long axis of the large cardioid being $d = 10$) and hybrid LDDMM with the same kernel and $\langle \cdot, \cdot \rangle_q$ given by (9). Both approaches perfectly align the template to the target, but their solutions differ. The LDDMM trajectories exhibit a typical behavior in which points tend to space out during motion; this behavior is not observed in the hybrid LDDMM trajectories, because (9) penalizes changes of parametrization. This can be seen in Figure 3, in which the deforming template is plotted in red along a geodesic path, with green dots marking the discretized points (the same color code being used in subsequent figures). The difference between the estimated registrations can also be appreciated in the last two panels of Figure 2. Here, and in the following experiments, we used $\lambda = 1$ in $\|v\|_q$, and

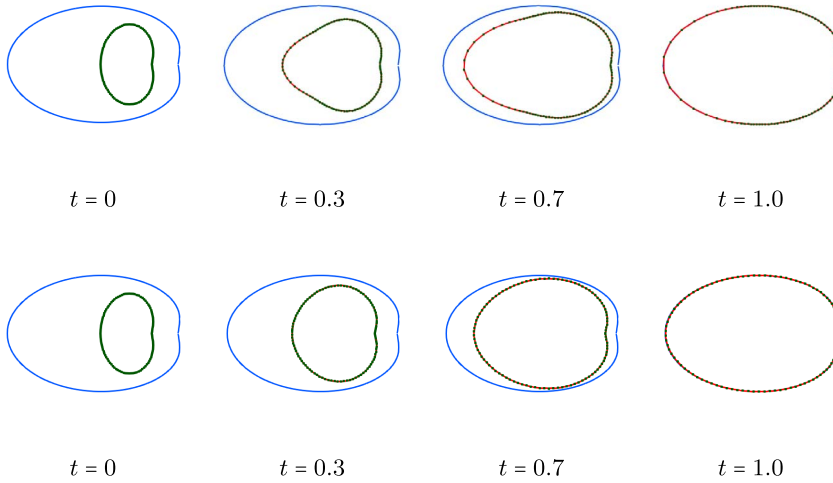


Figure 3: Smoothed Cardioids: geodesics. First row: LDDMM. Second row: Hybrid LDDMM. The initial discretization of the template is uniform. It remains close to uniform along the hybrid LDDMM geodesic path, while point are spacing out in the left side and accumulating on the right side along LDDMM geodesics.

added a multiplicative factor (between 200 and 500) in front of $\|\cdot\|_q$ when running the hybrid version.

2.3.2. Nested ellipses. Our second example is more challenging and involves multiple curves. Both template and target are composed with two small ellipses included in a large one (see Figure 4). For registration, the large ellipses are paired with each other, while the small ellipses are switched, i.e., the one on the left in the template is paired with the one on the right in the target and vice versa. This is achieved by defining an end-point term as

$$D(q^{\text{large}}, q_1^{\text{large}}) + D(q^{\text{left}}, q_1^{\text{right}}) + D(q^{\text{right}}, q_1^{\text{left}})$$

where D is given by (14).

The geodesics estimated with each method differ significantly and show interesting features. With standard LDDMM, we keep observing large reparametrization of each of the three curves, similar to what we observed in the previous example. The small ellipses avoid each other when changing places by flattening their shapes. We ran Hybrid LDDMM with $\|\cdot\|_q$ given by (9) and (8). In both cases, the reparametrization is uniform along




Figure 4: Nested Ellipses: Estimated deformations. (a) Template (red) and target (blue); (b) Standard LDDMM; (c) Hybrid LDDMM with (9). (d) Hybrid LDDMM with (8) Note that grid lines in the latter case are crossing over in this last panel. This is due to the resolution of the discretization used for this illustration compared to the size of the estimated deformation, which is nonetheless diffeomorphic, as implied by the curves remaining non-intersecting in the geodesic path.

each curve. With (9), which is scale and rotation invariant, the small ellipses shrink when crossing each other, before growing back to match the target. When using (8) (which is only rotation invariant), shrinking is not free anymore, and the curves make a wide berth to avoid each other. The kernel width was the same in all three experiments, in which we took $a = 0.2$.

2.3.3. Rays. We now compare configurations of $m = 10$ line segments stemming from a common origin (Figure 6). The segments' orientations are sampled uniformly over $[0, 2\pi]$ ($\theta_k = 2k\pi/m$, $k = 0, \dots, m - 1$) in the template, but not in the target ($\theta_k = 2\pi\sqrt{k/m}$, $k = 0, \dots, m - 1$). The target is moreover slightly translated. Here, and for the examples that follow, the cost function considers the curves as unlabeled (no correspondence information is used). Formally, this corresponds to considering that the curves are parametrized over the unions of m copies of M , $M^{(m)} = \bigcup_{k=1}^m (\{k\} \times M)$ and using $M^{(m)}$ in place of M in (15). This choice makes the matching problem significantly harder, creating possible local minima in the cost function.

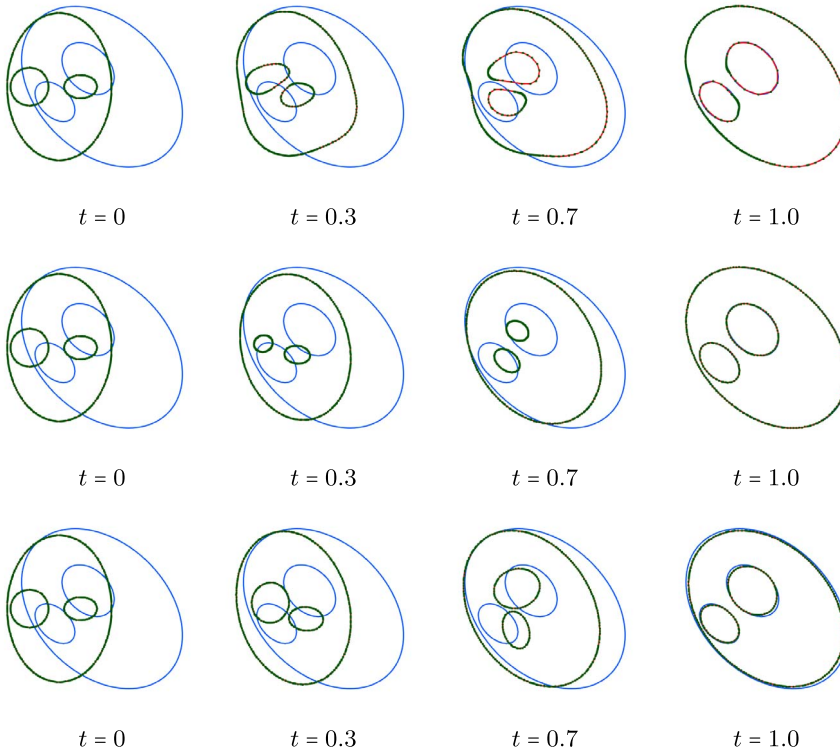


Figure 5: Ellipses: geodesics. First row: LDDMM. Second row: Hybrid LDDMM. The initial discretization of the template is uniform.

Such local minima actually trap the LDDMM algorithm when using small kernel sizes, and the solution provided in our experiments use a rather large kernel size, $a = L/5$, where L is the common lengths of the segments. The hybrid model uses $a = L/25$ combined with (8), the H^1 norm corrected for rotations.

As a result of the use of a large kernel in the LDDMM case, the obtained solution does not achieve a perfect transformation of the first segment (the one requiring the largest rotation) which is curved at the endpoint of the geodesic (see Figure 7). The segments remain perfectly straight along the geodesic estimated with the hybrid norm (visually at least: an exact transformation of the rays would not be diffeomorphic, but the deviation from a straight line happens below the discretization level chosen for the curves). The effect of the kernel size is also apparent in the estimated transformations, depicted in Figure 6. Similar to the previous examples,




Figure 6: Rays: Estimated deformations. (a) Template (red) and target (blue); (b) Standard LDDMM; (c) Hybrid LDDMM.

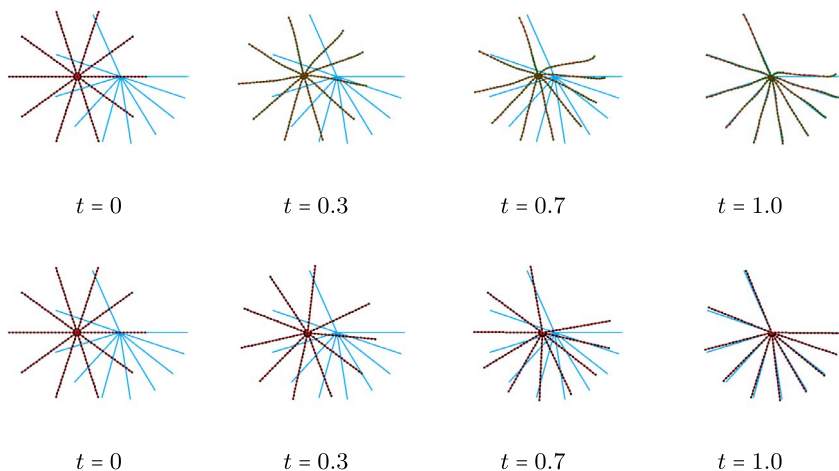


Figure 7: Rays: geodesics. First row: LDDMM. Second row: Hybrid LDDMM. The initial discretization of the template is uniform.

the reparametrization of the segments is more pronounced with standard LDDMM.

2.3.4. Half circles. Our last 2D example is similar to the previous one, using half circles, with various radii, instead of straight lines. We used the rotation- and scale-invariant norm (9) in the hybrid case, and the kernel sizes were $L/5$ and $L/25$ for standard and hybrid LDDMM, L being the radius of the largest circle. Both methods do a good job in registering the target to the template, but find different solutions as seen in Figure 8. The LDDMM solution tends to compress the space in the middle of the estimated pattern, while hybrid LDDMM estimates a motion closer to a rotation (which is cheap with the considered norm).




Figure 8: Arcs of circle: Estimated deformations. (a) Template (red) and target (blue); (b) Standard LDDMM; (c) Hybrid LDDMM.




Figure 9: Arcs of circle: geodesics. First row: LDDMM. Second row: Hybrid LDDMM. The initial discretization of the template is uniform.

3. Surfaces

The same approach can also be used with surfaces. At high level, not much needs to be modified formally from the curve case, simply letting M be the unit disc, or the unit sphere, or any other manifold sharing the topology of the considered surfaces. One can find a large collection of possible choices for $\{\cdot\}_q$ in [11, 8, 27, 10]. In the experiment that follows, we used one of the simplest options, letting

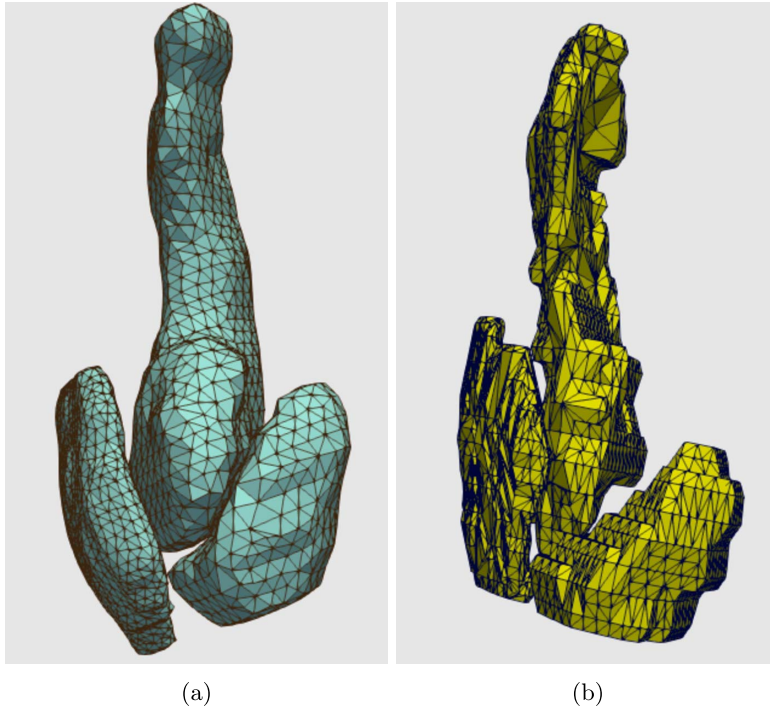


Figure 10: (a) Template and (b) target, each of them the union of three surfaces (left: entorhinal cortex, center: hippocampus, right: amygdala).

$$(16) \quad \|\!\|h\|\!\|_q^2 = \int_S |\nabla_S h|^2 d\sigma_S$$

where $S = q(M)$ with volume measure σ_S and Riemannian gradient ∇_S . (For this to be well defined, it is important that q remains an embedding at all times along finite energy paths.)

Our example uses the same data as the one presented in Figures 8 and 9 of [3]. It includes three shapes (see Figure 10) who are relatively close to each other (the hippocampus and the amygdala are actually slightly overlapping in the target). If one uses standard LDDMM with a small kernel ($a \simeq d/45$, where d is the height of the hippocampus), as illustrated in the first row of Figure 11, the diffeomorphism has undesirable properties, crunching parts of the surfaces (such as the front of the hippocampus, the bottom of the entorhinal cortex—which even has a residual spike—and the top of the amygdala) to match the target. With a larger kernel width ($a \simeq d/6$) the three shapes are transformed as if they formed one single




Figure 11: Surfaces: geodesics. First row: LDDMM (small kernel size). Second row: LDDMM (large kernel size). Third row: Hybrid LDDMM (small kernel size).

object, resulting in large reparametrization of the surfaces when they move along each other, because points that were nearby, initially tend to have similar motions. This is illustrated in the second row of Figure 11. The third row provides the geodesic obtained with the hybrid norm, with the same small kernel width as in the first row, but with (16) penalizing large deformations on the surfaces. In this case, the surfaces move nicely along each other, without requiring large reparametrization, except those required by the change in their respective shapes.

4. Conclusion

Our results illustrate several advantages of combining the standard LD-DMM approach with geometrically inspired norms on spaces of curves and surfaces. This very simple concept allows for much more modeling accuracy and flexibility, with a moderate computational impact. This is especially useful when dealing with complex configurations of shapes, as we saw in our examples.

There is clearly still room for future work and development, including the use of higher-order norms for curves and surfaces, and guidelines on which norm to use in specific applications. A version of the method for image matching is another important direction to be explored in the future. Formally, this requires defining $\|h\|_q$, where both q and h are scalar functions in \mathbb{R}^d and using the extra term $\|v^T \nabla q\|_q$ in the Riemannian norm. One option worth exploring is

$$\|h\|_q^2 = \int_{\mathbb{R}^d} |\nabla h(x)|^2 w_q(x) dx$$

where w_q is a “weight function” that depends on q , making, for example, deformations more costly in gray/white matter regions than within cerebrospinal fluid in brain mapping. This will be addressed in future work.

References

- [1] Sylvain Arguillère. The abstract setting for shape deformation analysis and lddmm methods. In Frank Nielsen and Frédéric Barbaresco, editors, *Geometric Science of Information: Second International Conference, GSI 2015, Palaiseau, France, October 28–30, 2015, Proceedings*, pages 159–167, Cham, 2015. Springer International Publishing. [MR3442196](#)

- [2] Sylvain Arguillère, Emmanuel Trélat, Alain Trouvé, and Laurent Younes. Shape deformation analysis from the optimal control viewpoint. *Journal de Mathématiques Pures et Appliquées*, 104(1):139–178, 2015. [MR3350723](#)
- [3] Sylvain Arguillère, Emmanuel Trélat, Alain Trouvé, and Laurent Younes. Registration of multiple shapes using constrained optimal control. *SIAM Journal on Imaging Sciences*, 9(1):344–385, 2016. [MR3472010](#)
- [4] J. Ashburner and K. J. Friston. Voxel based morphometry – the methods. *Neuroimage*, 11(6):805–821, 2000.
- [5] John Ashburner. A fast diffeomorphic image registration algorithm. *Neuroimage*, 38(1):95–113, 2007.
- [6] John Ashburner and Karl J. Friston. Diffeomorphic registration using geodesic shooting and gauss–newton optimisation. *NeuroImage*, 55(3):954–967, 2011.
- [7] Brian Avants and James C. Gee. Geodesic estimation for large deformation anatomical shape averaging and interpolation. *Neuroimage*, 23:S139–S150, 2004.
- [8] Martin Bauer and Martins Bruveris. A new Riemannian setting for surface registration setting for surface registration. In *Proceedings of the Third International Workshop on Mathematical Foundations of Computational Anatomy-Geometrical and Statistical Methods for Modelling Biological Shape Variability*, pages 182–193, 2011.
- [9] Martin Bauer, Martins Bruveris, Stephen Marsland, and Peter W. Michor. Constructing reparameterization invariant metrics on spaces of plane curves. *Differential Geometry and its Applications*, 34:139–165, 2014. [MR3209542](#)
- [10] Martin Bauer, Martins Bruveris, and Peter W. Michor. Overview of the geometries of shape spaces and diffeomorphism groups. *Journal of Mathematical Imaging and Vision*, 50(1–2):60–97, 2014. [MR3233135](#)
- [11] Martin Bauer, Philipp Harms, and Peter W. Michor. Sobolev metrics on shape space of surfaces. *Journal of Geometric Mechanics*, 389–438, 2011. [MR2888014](#)
- [12] M. Faisal Beg, Michael I. Miller, Alain Trouvé, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of

- diffeomorphisms. *International journal of computer vision*, 61(2):139–157, 2005.
- [13] Yan Cao, Michael I. Miller, Raimond L. Winslow, and Laurent Younes. Large deformation diffeomorphic metric mapping of vector fields. *Medical Imaging, IEEE Transactions on*, 24(9):1216–1230, 2005.
- [14] Can Ceritoglu, Kenichi Oishi, Xin Li, Ming-Chung Chou, Laurent Younes, Marilyn Albert, Constantine Lyketsos, Peter van Zijl, Michael I Miller, and Susumu Mori. Multi-contrast large deformation diffeomorphic metric mapping for diffusion tensor imaging. *Neuroimage*, 47(2):618–627, 2009.
- [15] N. Charon and A. Trouvé. The varifold representation of nonoriented shapes for diffeomorphic registration. *SIAM Journal on Imaging Sciences*, 6(4):2547–2580, 2013. [MR3138101](#)
- [16] Gary E. Christensen, Richard D. Rabbitt, and Michael I. Miller. Deformable templates using large deformation kinematics. *Image Processing, IEEE Transactions on*, 5(10):1435–1447, 1996.
- [17] Marc Droske and Martin Rumpf. A variational approach to nonrigid morphological image registration. *SIAM Journal on Applied Mathematics*, 64(2):668–687, 2004. [MR2049668](#)
- [18] Joan Glaunès, Anqi Qiu, Michael I. Miller, and Laurent Younes. Large deformation diffeomorphic metric curve mapping. *International journal of computer vision*, 80(3):317–336, 2008.
- [19] Monica Hernandez, Salvador Olmos, and Xavier Pennec. Comparing algorithms for diffeomorphic registration: Stationary lddmm and diffeomorphic demons. In *2nd MICCAI Workshop on Mathematical Foundations of Computational Anatomy*, pages 24–35, 2008.
- [20] Darryl D. Holm, Jerrold E. Marsden, and Tudor S. Ratiu. The Euler–Poincaré equations and semidirect products with applications to continuum theories. *Advances in Mathematics*, 137(1):1–81, 1998. [MR1627802](#)
- [21] Darryl D. Holm, Alain Trouvé, and Laurent Younes. The Euler–Poincaré theory of metamorphosis. *Quarterly of Applied Mathematics*, 97:661–685, 2009. [MR2588229](#)
- [22] S. Joshi. *Large Deformation Diffeomorphisms and Gaussian Random Fields for Statistical Characterization of Brain Sub-manifolds*. PhD thesis, Sever institute of technology, Washington University, 1997.

- [23] Sarang C. Joshi and Michael I. Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Transactions on Image Processing*, 9:1357–1370, 2000. [MR1808275](#)
- [24] E. Klassen, A. Srivastava, W. Mio, and S. H. Joshi. Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(3):372–383, 2004.
- [25] Arno Klein, Jesper Andersson, Babak A. Ardekani, John Ashburner, Brian Avants, Ming-Chang Chiang, Gary E. Christensen, D. Louis Collins, James Gee, Pierre Hellier, Joo Hyun Song, Mark Jenkinson, Claude Lepage, Daniel Rueckert, Paul Thompson, Tom Vercauteren, Roger P. Woods, J. John Mann, and Ramin V. Parsey. Evaluation of 14 nonlinear deformation algorithms applied to human brain {MRI} registration. *NeuroImage*, 46(3):786–802, 2009.
- [26] Hi L. Krall and Orrin Frink. A new class of orthogonal polynomials: The Bessel polynomials. *Transactions of the American Mathematical Society*, 65(1):100–115, 1949. [MR0028473](#)
- [27] Sebastian Kurtek, Eric Klassen, Zhaohua Ding, Sandra W. Jacobson, Joseph L. Jacobson, Malcolm J. Avison, and Anuj Srivastava. Parameterization-invariant shape comparisons of anatomical surfaces. *IEEE Transactions on Medical Imaging*, 30(3):849–858, 2011.
- [28] P. W. Michor and D. Mumford. Riemannian geometries on spaces of plane curves. *J. Eur. Math. Soc.*, 8:1–48, 2006. [MR2201275](#)
- [29] P. W. Michor and D. Mumford. An overview of the Riemannian metrics on spaces of curves using the Hamiltonian approach. *Applied and Computational Harmonic Analysis*, 23(1):74–113, 2007. [MR2333829](#)
- [30] Michael I. Miller, Alain Trounev, and Laurent Younes. Geodesic shooting for computational anatomy. *Journal of mathematical imaging and vision*, 24(2):209–228, 2006. [MR2227097](#)
- [31] Michael I. Miller and Laurent Younes. Group actions, homeomorphisms, and matching: A general framework. *International Journal of Computer Vision*, 41(1–2):61–84, 2001.
- [32] Dimitrios Pantazis, Richard M. Leahy, Thomas E. Nichols, and Martin Styner. Statistical surface-based morphometry using a nonparametric approach. In *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*, pages 1283–1286. IEEE, 2004.

- [33] Laurent Risser, F. Vialard, Robin Wolz, Maria Murgasova, Darryl D. Holm, and Daniel Rueckert. Simultaneous multi-scale registration using large deformation diffeomorphic metric mapping. *Medical Imaging, IEEE Transactions on*, 30(10):1746–1759, 2011.
- [34] Martin Styner, Ipek Oguz, Shun Xu, Christian Brechbühler, Dimitrios Pantazis, James J Levitt, Martha E Shenton, and Guido Gerig. Framework for the statistical shape analysis of brain structures using spharm-pdm. *The insight journal*, 1071:242, 2006.
- [35] Alain Trouvé and Laurent Younes. Metamorphoses through Lie group action. *Foundations of Computational Mathematics*, 5(2):173–198, 2005. [MR2149415](#)
- [36] Tom Vercauteren, Xavier Pennec, Aymeric Perchant, and Nicholas Ayache. Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage*, 45(1):S61–S72, 2009.
- [37] Anthony Yezzi and Andrea Mennucci. Metrics in the space of curves. *arXiv preprint math/0412454*, 2004.
- [38] L. Younes, P. Michor, J. Shah, and D. Mumford. A metric on shape spaces with explicit geodesics. *Rend. Lincei Mat. Appl.*, 9:25–57, 2008. [MR2383560](#)
- [39] Laurent Younes. Computable elastic distances between shapes. *SIAM Journal on Applied Mathematics*, 58(2):565–586, 1998. [MR1617630](#)
- [40] Laurent Younes. Jacobi fields in groups of diffeomorphisms and applications. *Quarterly of applied mathematics*, 65(1):113–134, 2007. [MR2313151](#)
- [41] Laurent Younes, Felipe Arrate, and Michael I. Miller. Evolutions equations in computational anatomy. *NeuroImage*, 45(1):S40–S50, 2009.

LAURENT YOUNES

DEPARTMENT OF APPLIED MATHEMATICS AND STATISTICS

AND CENTER FOR IMAGING SCIENCE

JOHNS HOPKINS UNIVERSITY

E-mail address: laurent.younes@jhu.edu

RECEIVED JUNE 5, 2017