Diffeomorphic Registration with Density Changes for the Analysis of Imbalanced Shapes

Hsi-Wei Hsieh and Nicolas Charon(✉)

Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, MD 21218, USA
hhsieh9@jhu.edu, charon@cis.jhu.edu

Abstract. This paper introduces an extension of diffeomorphic registration to enable the morphological analysis of data structures with inherent density variations and imbalances. Building on the framework of Large Diffeomorphic Metric Matching (LDDMM) registration and measure representations of shapes, we propose to augment previous measure deformation approaches with an additional density (or mass) transformation process. We then derive a variational formulation for the joint estimation of optimal deformation and density change between two measures. Based on the obtained optimality conditions, we deduce a shooting algorithm to numerically estimate solutions and illustrate the practical interest of this model for several types of geometric data such as fiber bundles with inconsistent fiber densities or incomplete surfaces.

Keywords: Shape analysis · Diffeomorphic registration · Generalized measures · Density variations

1 Introduction

The field known as computational anatomy [7] has come a long way since its inception several decades ago. Being primarily focused on the development of computational and statistical tools for the analysis of anatomy and its variability, the discipline has thus thrived by benefiting on the one hand from the increasing availability of efficient imaging techniques that can generate large amount of anatomical data and on the other hand from the mathematical advances made in shape analysis that provide the adequate theoretical frameworks and numerical methods for morphological analysis. One of these important theoretical milestones came from the representation of shape spaces as infinite-dimensional Riemannian manifolds which provides the foundations not only to construct relevant families of metrics between geometric structures but also to extend many classical statistics and machine learning tools to those shape spaces.

H.-W. Hsieh and N. Charon—Both authors were supported by NSF grant 1945224.

A. Feragen et al. (Eds.): IPMI 2021, LNCS 12729, pp. 31–42, 2021.
https://doi.org/10.1007/978-3-030-78191-0_3

Among the different Riemannian shape analysis frameworks which have been introduced, deformation-based models [1,12] have found many applications to biomedical data. These are usually referred to as extrinsic (our outer) shape metrics since the distance between two shapes is induced by a metric on the deformation group and measured by how much deformation is needed to map one on the other, which involves solving a registration problem. In particular, the model known as Large Deformation Diffeomorphic Metric Mapping (LDDMM) [1] allows to formulate such a problem as finding an optimal flow map between the two shapes and has been successfully applied to objects such as landmarks, images as well as curves and surfaces. In the latter cases, one additional difficulty is the absence of predefined point to point correspondences between the vertices of the two shapes. Addressing this particular issue motivated new connections with the area of geometric measure theory beginning with the work of [5] and later pursued in e.g. [2–4,9]. These works all share the same guiding principle that shapes such as curves and surfaces are better viewed as elements in certain spaces of measures, which allows to build correspondence free divergences that can be nicely embedded in diffeomorphic registration formulations.

So far, those measure representations have rarely been exploited in shape analysis beyond the pure diffeomorphic transformation setting described above. However, there are important limitations to those methods when dealing with registration of what we term generically as *imbalanced shapes*, namely in the situation where the representing measures display significant variations of mass or density. A motivating example is the case of white matter fiber bundles in which one can expect not only variations in the overall geometry of the bundle but also changes in the number (i.e. density) of fiber curves in each bundle. Diffeomorphic registration of fiber bundles [6] thus typically rely on an ad hoc renormalization/simplification step to compensate for fiber density inconsistencies. Another quite common situation is when a shape, for instance an anatomical surface, is only partially or sparsely known due to acquisition or segmentation issues. We propose to take further advantage of the flexibility of the measure setting by augmenting the diffeomorphic component of LDDMM with a global or local change of density of the source measure to account for potential mass imbalance. We then introduce a generalized registration model in which deformation and density change are estimated jointly. Our model differs from the metamorphosis setting of [10] in that we consider a more general class of measures better adapted to curves and surfaces but also restrict to transformations of the density only, thus avoiding the singularity issues described in [10].

The paper is organized as follows. In Sect. 2, we review the necessary background on diffeomorphic registration and measure representation of shapes. Section 3 introduces our generalized model for diffeomorphic registration with density variations as well as the proposed optimization algorithm. Numerical results on real and synthetic data are presented in Sect. 4.

2 Diffeomorphic Registration of Geometric Measures

The model that we develop in Sect. 3 draws from the concepts of diffeomorphic flows for the modelling of deformations on the one hand and of generalized

measures for the representation of shapes on the other. In this section, we give a brief overview of these two mathematical building blocks underlying our paper.

2.1 Diffeomorphisms and Registration

The construction of deformations as flows of vector fields goes back to works such as [1] and [12]. In this paper, we adopt the LDDMM setting of [1]. Let $n \in \mathbb{N}$ with $n \geq 2$ and $\mathbb{R}^n$ be the embedding space of the considered shapes. We will typically have $n = 2$ or $n = 3$ in our examples as we shall primarily be interested in 2D and 3D shapes. We will denote by V a Hilbert space of vector fields of $\mathbb{R}^n$ that is embedded in the space $C_0^1(\mathbb{R}^n, \mathbb{R}^n)$ of C^1 vector fields that vanish at infinity. Then let $L^2([0,1], V)$ be the space of time-varying vector fields such that for all $t \in [0,1]$, $v(t) \in V$ and $\int_0^1 \|v(t)\|_V^2 dt < +\infty$. The flow of $v \in L^2([0,1], V)$ at time $t \in [0,1]$ is defined as the mapping $\varphi_t^v : \mathbb{R}^n \to \mathbb{R}^n$ such that for all $x \in \mathbb{R}^n$, $\varphi_t^v(x) = x + \int_0^t v(s) \circ \varphi_s^v(x) ds$. It follows from the results of [14] (Chap. 7) that φ_t^v belongs to the set $\mathrm{Diff}^1(\mathbb{R}^n)$ of C^1-diffeomorphisms of $\mathbb{R}^n$, namely the bijective maps such that both φ_t^v and its inverse $(\varphi_t^v)^{-1}$ are C^1. Moreover, the term $\int_0^1 \|v(t)\|_V^2 dt$ where $\|\cdot\|_V$ is the norm of the Hilbert space V provides a measure of the energy of the deformation path $t \mapsto \varphi_t^v$ which is related to a right-invariant metric on the diffeomorphism group and is thus often used as a regularization energy for registration problems as we will see next. Furthermore, the embedding assumption on V implies that it is a Reproducing Kernel Hilbert Space (RKHS) of vector fields on $\mathbb{R}^n$. In all this paper, we shall make the additional (and quite common) assumption that the associated matrix kernel is of the form $K_V(x,y) = k_V(x,y)\mathrm{Id}_{n\times n}$ where $k_V : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ is a positive definite scalar kernel on $\mathbb{R}^n$.

Shapes can be then transformed by diffeomorphisms by specifying a group action: for instance, the action of a diffeomorphism $\phi \in \mathrm{Diff}^1(\mathbb{R}^n)$ on a set $q = (x_i)_{i=1,\ldots,N}$ of N points of $\mathbb{R}^n$ is typically defined by transporting each point by ϕ i.e. $\phi \cdot q = (\phi(x_i))_{i=1,\ldots,N}$. Registering two such point sets q and q' can be then formulated in this framework as the minimization over all $v \in L^2([0,1], V)$ of a functional like $\frac{1}{2}\int_0^1 \|v(t)\|_V^2 dt + \lambda \sum_{i=1}^N |\varphi_1^v(x_i) - x_i'|^2$ that is a weighted sum of the deformation energy and the squared distances between the corresponding points in $\phi \cdot q$ and q'. This is the problem known as (inexact) landmark registration.

2.2 Geometric Measure Representation of Shapes

In many problems of interest however, shapes cannot be directly described as landmarks. For instance, datasets of triangulated surfaces usually exhibit differences in sampling or do not come automatically with point correspondences across all the dataset. This has motivated the exploration of shape representations from geometric measure theory which allow the comparison and registration of geometric structures without the need for such correspondences. Measure frameworks such as currents or varifolds [5,9] provide a general setting to encode geometric shapes as unlabelled points in $\mathbb{R}^n$ that carry some information of local

tangent plane as well. In this paper, we will specifically consider as our shape space $\mathcal{S}$ the set of all objects represented by a discrete measure of the form $\mu = \sum_{i=1}^{N} r_i \delta_{(x_i, T_i)}$ for $N \geq 1$. In this representation, each Dirac $r_i \delta_{(x_i, T_i)}$ can be interpreted as an unlabelled particle of mass $r_i > 0$ (which we shall also refer to as the density of μ at x_i in a discrete sense) located at the position $x_i \in \mathbb{R}^n$ and carrying an oriented d-dimensional subspace T_i. Note that here $0 \leq d \leq n$ is fixed and we will be mostly interested in the cases $d = 1$ and $d = 2$ in practice, i.e. the T_i's are oriented lines or planes. Mathematically, μ is a positive measure on the product of $\mathbb{R}^n$ and the oriented d-dimensional Grassmannian of $\mathbb{R}^n$, and is usually called a d-current or an oriented d-varifold following the definitions and terminology of [8,9]. In practice, it will be more convenient to represent each oriented subspace T_i by an oriented frame of d linearly independent vectors $(u_i^{(k)})_{k=1,\ldots,d}$. Although such a frame is not unique, this will not constitute an issue for the applications considered in this work.

The above space $\mathcal{S}$ of discrete measures provides an effective setting to embed a variety of geometric structures. In particular, curves and surfaces can always be approximated by elements of $\mathcal{S}$ [3,9]. For instance a polygonal curve with edges $[e_i^1, e_i^2]$ for $i = 1, \ldots, N$ can be approximated by the measure of $\mathcal{S}$ (with $d = 1$) $\mu = \sum_{i=1}^{N} r_i \delta_{(x_i, T_i)}$ where r_i is the edge length, $x_i = (e_i^1 + e_i^2)/2$ its midpoint and T_i the oriented line directed by $\overrightarrow{e_i^1 e_i^2}$. Triangulated surfaces can be similarly approached by a 2-dimensional measure in which case r_i is the area of the face, x_i the barycenter of its three points and T_i the oriented plane containing the triangular face. Beyond curves and surfaces, $\mathcal{S}$ is a versatile class of objects represented fundamentally as a distribution of tangent spaces spread at different locations in $\mathbb{R}^n$. One of the key advantage of such measure representations of shapes is that one can easily equip $\mathcal{S}$ with a metric. Among the different classes of metrics between measures that one may choose from, kernel norms (also known as maximum mean discrepancy) are particularly well-suited in our context as they lead to relatively simple and explicit expression of the distance between two measures in $\mathcal{S}$. Specifically, given a positive-definite kernel $K_{\mathcal{S}}$ on the product of $\mathbb{R}^n$ and the set of all oriented d-planes of $\mathbb{R}^n$ (or more simply the set of all oriented d- frames), the associated Hilbert inner product between any $\mu = \sum_{i=1}^{N} r_i \delta_{(x_i, T_i)}$ and $\mu' = \sum_{j=1}^{N'} r'_j \delta_{(x'_j, T'_j)}$ in $\mathcal{S}$ is given by:

$$\langle \mu, \mu' \rangle_{\mathcal{S}} = \sum_{i=1}^{N} \sum_{j=1}^{N'} r_i r'_j K_{\mathcal{S}}(x_i, T_i, x'_j, T'_j). \tag{1}$$

Then the computation of the distance on $\mathcal{S}$ is just obtained from $\|\mu - \mu'\|_{\mathcal{S}}^2 = \|\mu\|_{\mathcal{S}}^2 - 2\langle \mu, \mu' \rangle_{\mathcal{S}} + \|\mu'\|_{\mathcal{S}}^2$ and reduces to evaluations of the kernel function. Note that this provides a notion of distance that does not rely on any correspondence between the Diracs of μ and μ' (and that actually remains well-defined even when the number of Diracs in μ and μ' are different). The properties of these metrics have been studied extensively in [5,8,9] and they provide a convenient notion of discrepancy on our space $\mathcal{S}$ for adequate choices of the kernel $K_{\mathcal{S}}$, such as the one we will specify and use in Sect. 4 for our simulations.

To arrive at a formulation of the LDDMM registration problem for measures in $\mathcal{S}$, the only missing element is the group action of $\mathrm{Diff}(\mathbb{R}^n)$ on $\mathcal{S}$. A standard action is by measure pushforward [5,8] which is given for any $\phi \in \mathrm{Diff}(\mathbb{R}^n)$ and any $\mu = \sum_{i=1}^N r_i \delta_{(x_i,T_i)} \in \mathcal{S}$ by

$$\phi \cdot \mu = \sum_{i=1}^{N} |J_{x_i}^{T_i}\phi| . r_i \ \delta_{(\phi(x_i), d_{x_i}\phi(T_i))} \in \mathcal{S} \tag{2}$$

where $d_{x_i}\phi$ denotes the differential of ϕ at x_i, $d_{x_i}\phi(T_i)$ the subspace spanned by the transported frame vectors $(d_{x_i}\phi(u_i^{(k)}))_{k=1,\dots,d}$ and $|J_{x_i}^{T_i}\phi|$ is the absolute value of the Jacobian determinant of ϕ along T_i at x_i given explicitly by $\sqrt{\det(d_{x_i}\phi(u_i^{(k)}) \cdot d_{x_i}\phi(u_i^{(l)}))_{k,l}}$. This group action happens to be consistent with the usual action of diffeomorphisms on d-dimensional submanifolds of $\mathbb{R}^n$ in the sense that if $\mu \in \mathcal{S}$ is a discrete approximation of a submanifold M then $\phi \cdot \mu$ is typically an approximation of the submanifold $\phi(M)$, c.f. [8] for rigorous statements. This is the main reason for the presence of the mass changes $|J_{x_i}^{T_i}\phi|$ in (2) that represent the local change in d-volume induced by ϕ.

Finally the registration of a given source measure $\mu_0 \in \mathcal{S}$ to a target $\mu' \in \mathcal{S}$ can be framed as minimizing the functional $\frac{1}{2}\int_0^1 \|v(t)\|_V^2 dt + \lambda \|\varphi_1^v \cdot \mu_0 - \mu'\|_{\mathcal{S}}^2$. This class of optimal control problems includes as particular cases the curve and surface registration approaches of [2,5]. We will elaborate on the numerical aspects for the more general framework that we present in the following section.

3 Diffeomorphic Registration with Density Changes

Although the action of diffeomorphisms on $\mathcal{S}$ does allow to transform the mass of measures through the Jacobian determinant of the deformation, registering measures with important inconsistencies or density variations may lead to unnatural or even degenerate optimal deformations as we will show in some of the examples of Sect. 4. Our goal is thus to augment (2) with a complementary process to simultaneously modify the density of the measure. The main focus will be on a global model with a single common density rescaling factor but we will also briefly introduce a preliminary extension of the approach to deal with local changes in density as well.

3.1 An Augmented Optimal Control Problem

Adopting the notations of the previous section, we introduce a complementary rescaling factor $\alpha \in \mathbb{R}^+$, which is a nonnegative number acting as a global multiplicative factor on the measure μ which we write $\alpha.\mu$. Under this extended setting, we formulate the registration of a source μ_0 to a target μ' as the following new optimization problem:

$$\min_{v,\alpha} E(v,\alpha) \doteq \frac{1}{2}\int_0^1 \|v(t)\|_V^2 dt + \frac{\tau}{2}(\alpha-1)^2 + \lambda \|\alpha.\mu(1) - \mu'\|_{\mathcal{S}}^2 \tag{3}$$

subject to $\mu(t) \doteq (\varphi_t^v) \cdot \mu_0$. The rescaling factor α is here penalized by the simple squared difference with $\alpha = 1$ weighted by a fixed coefficient $\tau > 0$ and one can see formally that letting $\tau \to +\infty$ imposes $\alpha = 1$ and (3) then reduces to the previous LDDMM registration problem.

From now on, let (v, α) be a minimizer of (3) and $q(t)$ the associated optimal trajectory. We shall derive some necessary conditions satisfied by such a minimizer. We write $\mu_0 = \sum_{i=1}^N r_i \delta_{(x_i, T_i)}$ and, using the representation of the subspaces T_i by frames of d vectors as explained above, we can alternatively view the state variable of the optimal control problem as $q = ((x_i, u_i^{(k)})_{1 \leq i \leq N, 1 \leq k \leq d})$, where $\text{Span}(u_i^{(1)}, \ldots, u_i^{(d)}) = T_i$ and $|u_i^{(1)} \wedge \cdots \wedge u_i^{(d)}| = \sqrt{\det(u_i^{(k)} \cdot u_i^{(l)})} = r_i$.

We first notice that, as a function of α with v and $\mu(t)$ being fixed, E is quadratic and solving for $\frac{\partial E}{\partial \alpha} = 0$ shows that the optimal α can be expressed with respect to the final measure $\mu(1)$ as:

$$\alpha^* = \frac{\frac{\tau}{2} + \lambda \langle \mu(1), \mu' \rangle_{\mathcal{S}}}{\frac{\tau}{2} + \lambda \|\mu(1)\|_{\mathcal{S}}^2}. \tag{4}$$

where the Hilbert product on $\mathcal{S}$ is given by (1). Inserting into (3), this now allows to reduce the problem to an optimal control problem with control v. The optimality conditions on v can be derived similarly to [8,11,13], by introducing the Hamiltonian of the problem which is given by:

$$H(p, q, v) \doteq \sum_{i=1}^N \langle p_i^x, v(x_i) \rangle + \sum_{i=1}^N \sum_{k=1}^d \langle p_i^{u_k}, d_{x_i} v(u_i^{(k)}) \rangle - \frac{1}{2} \|v\|_V^2,$$

where $p_i^x, p_i^{u_k} \in \mathbb{R}^n$ denote the costates of the position x and frame vectors $u_i^{(k)}$. By applying the Pontryagin maximum principle, we find that any optimal trajectory $(x_i(t), u_i^{(k)}(t))$ must satisfy the following Hamiltonian equations:

$$\begin{cases} \dot{x}_i(t) = v(t)(x_i(t)) \\ \dot{u}_i^{(k)}(t) = d_{x_i(t)} v(t)(u_i^{(k)}(t)) \\ \dot{p}_i^x(t) = -d_{x_i(t)} v(t)^T p_i^x(t) - \sum_{k=1}^d d^{(2)}_{x_i(t)} v(t)(\cdot, u_i^{(k)}(t))^T p_i^{u_k}(t) \\ \dot{p}_i^{u_k}(t) = -d_{x_i(t)} v(t)^T p_i^{u_k}(t) \end{cases} \tag{5}$$

and, using the RKHS property of V, the optimal control v is given by:

$$v(t)(\cdot) = \sum_{i=1}^N k_V(x_i(t), \cdot) p_i^x(t) + \sum_{k=1}^d \partial_1 k_V(x_i(t), \cdot)(u_i^{(k)}(t), p_i^{u_k}(t)). \tag{6}$$

Thus, from the above equations, we obtain that the full energy functional to be minimized can be written as a function of the initial costates, namely

$$\begin{aligned} E(p_i^x(0), p_i^{u_k}(0)) = &\frac{1}{2} \sum_{i=1}^N \langle p_i^x(0), v(0)(x_i(0)) \rangle + \frac{1}{2} \sum_{i=1}^N \sum_{k=1}^d \langle p_i^{u_k}(0), d_{x_i(0)} v(0)(u_i^{(k)}(0)) \rangle \\ &+ \frac{\tau}{2} (\alpha^* - 1)^2 + \lambda \|\alpha^* . \mu(1) - \mu'\|_{\mathcal{S}}^2, \end{aligned} \tag{7}$$

where $\alpha^* . \mu(1) = \sum_{i=1}^N \alpha^* r_i(1) \delta_{(x_i(1), T_i(1))}$ and $T_i(1) = \text{Span}(\{u_i^{(k)}(1)\})$ is obtained from the Hamiltonian equations (5) and (6).

3.2 Numerical Implementation

The numerical minimization of the energy (7) can be tackled based on an iterative shooting scheme similar to other LDDMM approaches [8,13]. Specifically, given the initial costates $(p_i^x(0), (p_i^{u_k}(0))_{k=1}^d)$ at the current iteration of the algorithm together with the known and fixed initial state variables $(x_i(0), (u_i^{(k)}(0))_{k=1}^d)$, we start by integrating the Hamiltonian equations (5) and (6) based on an RK4 scheme to obtain the measure $\mu(1)$ at the final time. We then compute α^* with (4) from which we obtain the value of the energy (7). In order to update the initial costates, we also need the gradient of E which we can directly compute using automatic differentiation. More precisely, our Python implementation leverages the Pytorch library together with the recently developed KeOps library[1]. The latter allows to generate efficient CUDA subroutines for the computation and automatic differentiation of expressions involving positive definite kernels such as the ones appearing in the Hamiltonian equations and in the inner product of $\mathcal{S}$ given by (1). Finally, with E and ∇E being obtained as just explained, the optimization itself is done using the L-BFGS algorithm of the SciPy library.

In what follows, we will refer to this registration algorithm with global density rescaling by the acronym LDDMM+GD. The parameters that need to be set by the user are the kernels k_V and $K_{\mathcal{S}}$ as well as the weighing coefficients λ and τ. The latter controls the relative importance of deformation and mass rescaling in the overall change of density. We illustrate the effect of τ on the simplest example of two single Diracs in Fig. 1 (with $n = 2$ and $d = 1$). The optimal diffeomorphism φ_1^v pictured here via the resulting deformed grid shows a combination of a local rotation effect (in order to match the directions of the frame vectors) and of a local compression (to compensate for the difference in mass). The case $\tau = \infty$ corresponds to the pure diffeomorphic registration setting of the previous section. In sharp contrast, when $\tau = 0$, the deformation reduces to only rotating the directional component of the source Dirac while the transformation of mass is entirely done by the rescaling variable α^*. Intermediate values of τ lead to both φ_1^v and α^* contributing to the change in density.

3.3 Local Density Changes

The model presented in the previous sections is well-suited when a common and global density rescaling effect is expected as the results of Sect. 4 will illustrate but is typically not adapted to the situation of local mass imbalances such as in the case of particular missing parts on the target shape. To tackle this more general case, we briefly discuss a preliminary approach that can be derived as a localized version of the above model. Instead of the single density rescaling variable α, one can introduce N controls $\alpha_i \in \mathbb{R}_+$ associated to each Dirac in μ_0. Writing now $\alpha = (\alpha_1, \ldots, \alpha_N)$ and defining the measure mass rescaling as

[1] https://www.kernel-operations.io/.

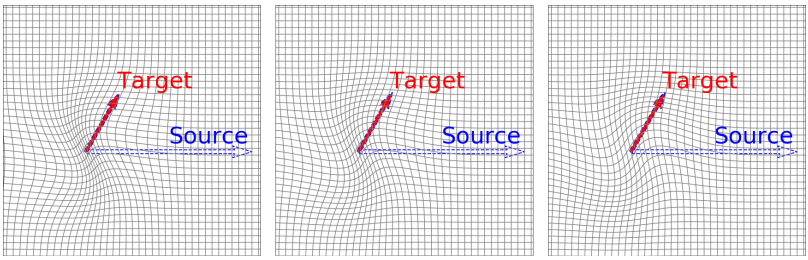

(a) $\tau = \infty$ (b) $\tau = 50$ (c) $\tau = 0$

Fig. 1. Registration between the two Diracs $\mu = r_0\delta_{(x_0,T_0)}$ and $\mu' = r'\delta_{(x',T')}$ with $r_0 = 2$, $r' = 4/5$, $x_0 = x' = (0,0) \in \mathbb{R}^2$ and T, T' are the lines spanned by $u_0 = (1,0)$ and $u' = (\cos(\pi/3), \sin(\pi/3))$. The plotted arrows vectors here represent r_0u_0 and $r'u'$ respectively. The figure illustrates the effect of the choice of τ on the registration, with $\tau = \infty$ corresponding to the pure deformation case (i.e. $\alpha^* = 1$). The optimal density rescaling factors in (b) and (c) are $\alpha^* = 0.6773$ and $\alpha^* = 0.4074$ respectively.

$\alpha \cdot \mu = \sum_{i=1}^N \alpha_i r_i \delta_{(x_i,T_i)}$, we formulate the registration problem with local density changes as follows:

$$\min_{v,\alpha\in\mathbb{R}_+^N} E(v,\alpha) \doteq \frac{1}{2}\int_0^1 \|v(t)\|_V^2 dt + \frac{\tau}{2}\sum_{i=1}^N r_i(\alpha_i - 1)^2 + \lambda\|\alpha\cdot\mu(1) - \mu'\|_{\mathcal{S}}^2 \quad (8)$$

in which the penalty on the density rescaling vector α is now the distance to $(1,\ldots,1)$ for the L^2 metric weighed by the density of the initial μ_0. Note that this is only one simple choice of penalty that directly extends (3) but we plan, in future work, to examine other relevant regularizers that would for instance constrain the local variations of the α_i's.

Problem (8) can be solved in similar fashion as in Sect. 3.2. Indeed, the Hamiltonian equations (5) and (6) still hold although the optimality equations on the α_i's are not as straightforward to exploit as (4). In practice, we instead jointly optimize over the initial costates $(p_i^x(0), (p_i^{u_k}(0))_{k=1}^d)$ together with $\alpha = (\alpha_i)$ using L-BFGS, with the gradient of the energy with respect to each α_i being computed by automatic differentiation. We will denote this diffeomorphic registration under local density changes by the acronym LDDMM+LD.

4 Results

As proof-of-concept, we present a few results of the above LDDMM+GD and LDDMM+LD algorithms applied to 3D shapes ($n = 3$), specifically discrete curves ($d = 1$) or surfaces ($d = 2$). Those shapes are converted to elements of $\mathcal{S}$ as explained in Sect. 2.2. For the purpose of visualization however, we shall plot the shapes rather than their associated measures in $\mathcal{S}$ and display the source shape's time evolution along the estimated deformation path φ_t^v. In all experiments, the

deformation kernel k_V is a Gaussian and the kernel $K_{\mathcal{S}}$ defining the metric on $\mathcal{S}$ is chosen among the class of kernels discussed in [9], specifically as the tensor product of a Gaussian kernel on $\mathbb{R}^n$ and the Binet kernel between d-dimensional subspaces, i.e. for $T = \text{Span}(u^{(k)})$ and $T' = \text{Span}(u'^{(k)})$ the positive kernel given by $\det((u^{(k)} \cdot u'^{(l)})_{k,l})^2$. We also set τ to a small value so as to put only minimal constraints on the estimation of α.

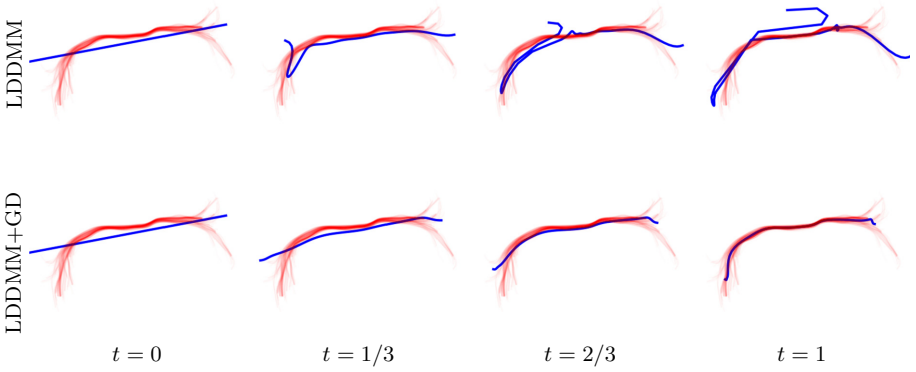

Fig. 2. Registration of single curve to CP fiber bundle (365 curves). The second row shows the deformation at intermediate times for the proposed LDDMM+GD algorithm where the estimated density rescaling is $\alpha^* = 348.80$.

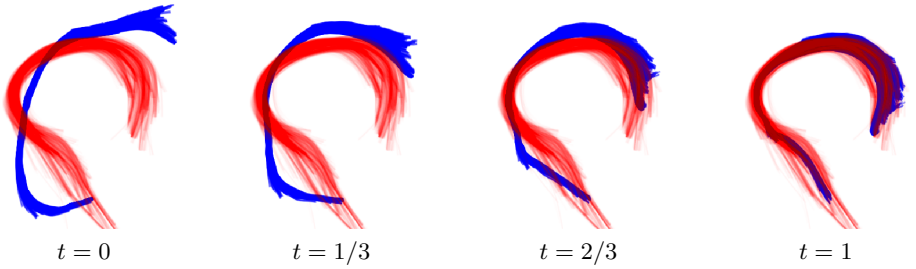

Fig. 3. Registration of CA fiber bundle (431 curves) to Fornix fiber bundle (2151 curves). The estimated α^* is 4.18, close to the fiber density ratio of the bundles.

Fiber Bundles. For our first set of simulations, we consider white matter fiber tracts taken from the publicly available ISMRM 2015 Tractography Challenge repository[2]. In Fig. 2, we show the result of registering a single template curve onto the posterior commissure (CP) bundle containing 365 distinct curves. As can be seen on the first row, regular LDDMM registration generates a folding of the source curve in an attempt to compensate for the difference in total mass.

[2] http://www.tractometer.org/ismrm_2015_challenge/.

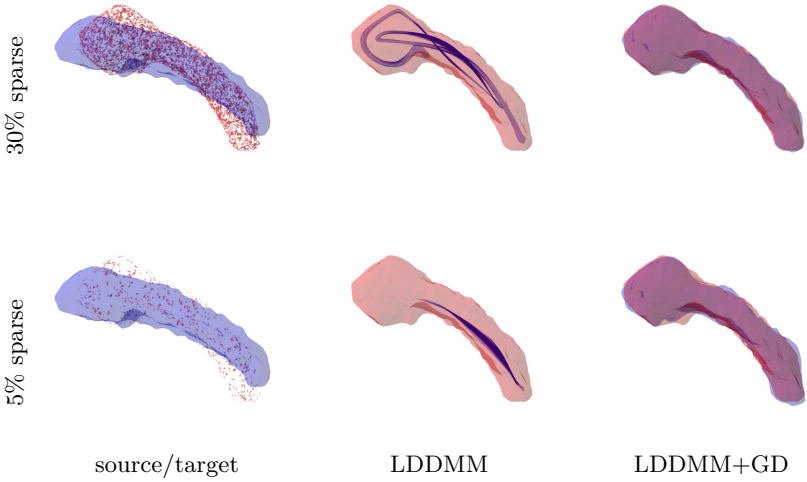

Fig. 4. Registration between two hippocampus surfaces: source (in blue) and target (in red) which has been randomly subsampled to 30% (first row) and 5% (second row) of its total number of triangles. The resulting deformed surface at $t = 1$ obtained with LDDMM (second column) and LDDMM+GD (third column) is compared to the fully sampled target surface. For LDDMM+GD, we obtain $\alpha^* = 0.29$ and $\alpha^* = 0.049$. (Color figure online)

The LDDMM+GD algorithm on the other hand leads to a deformed curve that matches the average geometry of the bundle with an estimated $\alpha^* = 348.80$ consistent with the density of curves in the target. Note that α^* is in fact smaller than 365 which accounts for the spatial spreading and fanning of the bundle. We also consider the registration between two different fiber tracts: the anterior commissure (CA) and fornix which are made of respectively 431 and 2151 individual curves. Once again, standard diffeomorphic registration (not shown here for the sake of space) induces important artifactual folding effects in contrast with the LDDMM+GD registration result of Fig. 3.

Sparse Shapes. Another feature of the LDDMM+GD approach is its robustness under sparse and incomplete observation of the target surface. This is illustrated on the example of hippocampi surfaces (data provided with the Keops library) in Fig. 4 where sparse targets are synthetically generated by keeping only a small number of random faces from the full ground truth target mesh. This mass imbalance results in severe shrinking and twisting of the registered surface estimated with standard LDDMM while LDDMM+GD recovers a surface close to the ground truth and automatically estimates (through α^*) the sparsity rate with good accuracy.

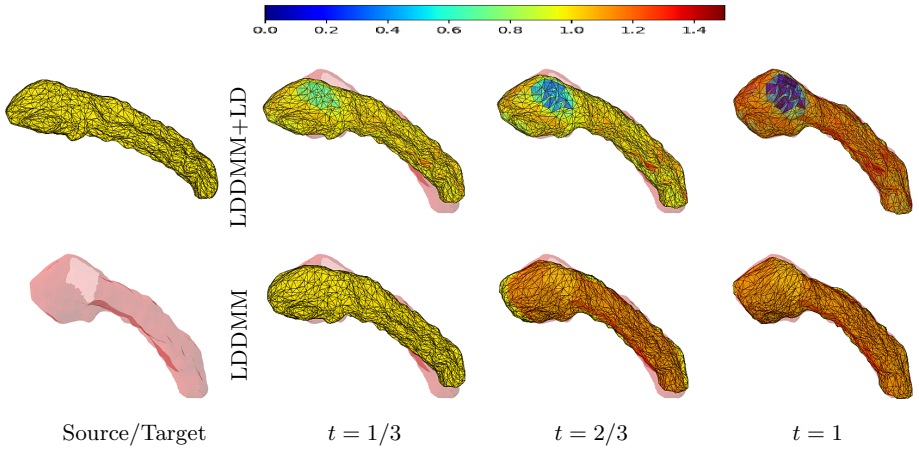

Fig. 5. Registration of two hippocampi surfaces with a missing subregion obtained with the LDDMM+LD and standard LDDMM approach. The colors on the first row correspond to the values of the mass rescaling α_i at each location. One can notice small differences in terms of overlap between the final shape and the ground truth target around that subregion. The Hausdorff distance to the ground truth is 1.5714 (LDDMM+LD) versus 1.9789 (LDDMM). (Color figure online)

Partial Matching. This approach is however not well-suited in the situation where an isolated part of the target shape is locally missing, as with the simulated example of Fig. 5 in which we artificially remove a subregion of the hippocampus surface. In such a case, the LDDMM+LD algorithm allows for local mass changes and is able to estimate, alongside the deformation, the corresponding missing region on the source shape as shown on the first row, where the colors represent the values of the α_i's at the different locations. In contrast, such missing regions can adversely affect registration under the standard LDDMM model. This is evidenced quantitatively by the closer proximity (measured for the usual Hausdorff distance) of the matched surface to the ground truth (i.e. complete) target for our proposed LDDMM+LD approach.

5 Conclusion

We introduced novel frameworks and algorithms for the registration of shapes modelled as discrete measures in which the deformation is coupled with a global or local transformation of the density. The Python implementation will be made openly available in the future and can be currently shared upon request. Our preliminary experiments hint at the potential of those models for the morphological analysis of partial shapes, a recurring issue of deformation models in computational anatomy, and for applications to shape completion. It is also well suited for registration or template estimation of fiber bundle data and could be used

in conjunction with other numerical methods in the field, such as the streamline approximation schemes of [6]. Future work will focus on pushing further those applications and on a more thorough analysis of the LDDMM+LD model, by investigating other possible regularizers on the density change and exploring connections with related unbalanced frameworks in optimal transport [4].

References

1. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. Int. J. Comput. Vis. **61**, 139–157 (2005)
2. Charon, N., Trouvé, A.: The varifold representation of nonoriented shapes for diffeomorphic registration. SIAM J. Imaging Sci. **6**(4), 2547–2580 (2013)
3. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N.: Statistical models of sets of curves and surfaces based on currents. Med. Image Anal. **13**(5), 793–808 (2009)
4. Feydy, J., Charlier, B., Vialard, F.-X., Peyré, G.: Optimal transport for diffeomorphic registration. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017, Part I. LNCS, vol. 10433, pp. 291–299. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66182-7_34
5. Vaillant, M., Glaunès, J.: Surface matching via currents. In: Christensen, G.E., Sonka, M. (eds.) IPMI 2005. LNCS, vol. 3565, pp. 381–392. Springer, Heidelberg (2005). https://doi.org/10.1007/11505730_32
6. Gori, P., et al.: Parsimonious approximation of streamline trajectories in white matter fiber bundles. IEEE Trans Med. Imaging **35**(12), 2609–2619 (2016)
7. Grenander, U., Miller, M.I.: Computational anatomy: an emerging discipline. Q. Appl. Math. **I**(4), 617–694 (1998)
8. Hsieh, H.-W., Charon, N.: Metrics, quantization and registration in varifold spaces. arXiv preprint arXiv:1903.11196 (2019)
9. Kaltenmark, I., Charlier, B., Charon, N.: A general framework for curve and surface comparison and registration with oriented varifolds. In: Computer Vision and Pattern Recognition (CVPR) (2017)
10. Richardson, C., Younes, L.: Computing metamorphoses between discrete measures. J. Geom. Mech. **5**(1), 131 (2013)
11. Sommer, S., Nielsen, M., Darkner, S., Pennec, X.: Higher-order momentum distributions and locally affine LDDMM registration. SIAM J. Imaging Sci. **6**(1), 341–367 (2013)
12. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: efficient non-parametric image registration. NeuroImage **45**(1), 61–72 (2009)
13. Vialard, F.-X., Risser, L., Rueckert, D., Cotter, C.: Diffeomorphic 3D image registration via geodesic shooting using an efficient adjoint calculation. Int. J. Comput. Vis. **97**(2), 229–241 (2012). https://doi.org/10.1007/s11263-011-0481-8
14. Younes, L.: Shapes and Diffeomorphisms. Springer, Heidelberg (2019). https://doi.org/10.1007/978-3-662-58496-5