PDiff: Irrotational Diffeomorphisms for Computational Anatomy

Jacob Hinkle and Sarang Joshi

Scientific Computing and Imaging Institute, Department of Bioengineering,
University of Utah, Salt Lake City, UT.

Abstract. The study of diffeomorphism groups is fundamental to computational anatomy, and in particular to image registration. One of the most developed frameworks employs a Riemannian-geometric approach using right-invariant Sobolev metrics. To date, the computation of the Riemannian log and exponential maps on the diffeomorphism group have been defined implicitly via an infinite-dimensional optimization problem. In this paper we the employ Brenier's (1991) polar factorization to decompose a diffeomorphism h as $h(x) = S \circ \psi(x)$, where $\psi = \nabla\rho$ is the gradient of a convex function ρ and $S \in \mathrm{SDiff}(\mathbb{R}^d)$ is a volume-preserving diffeomorphism. We show that all such mappings ψ form a submanifold, which we term $\mathrm{PDiff}(\mathbb{R}^d)$, generated by irrotational flows from the identity. Using the natural metric, the manifold $\mathrm{PDiff}(\mathbb{R}^d)$ is flat. This allows us to calculate the Riemannian log map on this submanifold of diffeomorphisms in closed form, and develop extremely efficient metric-based image registration algorithms. This result has far-reaching implications in terms of the statistical analysis of anatomical variability within the framework of computational anatomy.

Keywords: image registration, computational anatomy, irrotational, Helmholtz-Hodge decomposition, polar factorization

1 Introduction

Over the last decade, the field of computational anatomy has substantially matured and several approaches have been developed for the study of anatomical variations that are evident within medical images. The most theoretically developed and principled approaches are based on the Riemannian geometry of groups of diffeomorphisms of three-dimensional Euclidian space, $\mathbb{R}^3$, and its submanifolds (points, curves and surfaces) on which these groups act. Fundamental to this approach is the computation of geodesics which provide *normal coordinates* via the Riemannian log and exponential maps allowing for statistical analysis of anatomical variability. Despite the elegance of the theory, universal adoption has been limited by the computational complexity of the resulting optimization problems, especially the need for infinite dimensional optimization to compute the geodesic and the log map. To mitigate the computational complexity, recently some [7] have suggested abandoning the intrinsic Riemannian geometric approach and taking an extrinsic Eulerian view of deformation based on stationary vector fields.

One of the major contributions of this paper is the use of the remarkable result by Brenier [3] concerning the polar factorization of diffeomorphisms (analogous to the polar factorization of matrices) to define a submanifold of irrotational diffeomorphisms which we call $\mathrm{PDiff}(\mathbb{R}^d)$. In this paper we show that this infinite-dimensional submanifold is generated by irrotational velocity fields and that furthermore using the natural metric this submanifold is flat, meaning that sectional curvature in every direction is 0. This theoretical result has far reaching consequences: for example, within this space the intrinsic or Fréchet mean is guaranteed to be unique. Another consequence of this remarkable result is that we are able to derive in closed form the Riemannian log map and compute the distance between any two diffeomorphisms within PDiff in closed form. In this paper we begin to explore the applications of this by developing extremely computationally efficient and numerically stable image registration algorithms.

2 Mathematical Background and Notation

Although diffeomorphisms in the context of image registration have been extensively studied for completeness we review the basic set up. A compactly supported diffeomorphism φ is a bijective map from $\mathbb{R}^d$ to $\mathbb{R}^d$ such that both φ and its inverse φ^{-1} are smooth and have compact support. The identity transformation id is a diffeomorphism as well as the composition of any two. As the inverse of a diffeomorphism is also a diffeomorphism, it implies that the set of all diffeomorphisms forms a group. The Lie algebra $\mathfrak{g}$ of the compactly supported diffeomorphism group of $\mathrm{Diff}(\mathbb{R}^d)$ consists of all compactly supported smooth vector fields on $\mathbb{R}^d$, equipped with the Lie bracket of vector fields.

Given a time dependent vector field $v(x,t)$ one defines a path in $\mathrm{Diff}(\mathbb{R}^d)$ via the O.D.E:

$$\frac{d\varphi(x,t)}{dt} = v(\varphi(x,t),t), \text{ with initial condition: } \varphi(x,0) = x.$$

One induces a right invariant metric by choosing a differential operator L which acts on velocity fields. This operator determines the norm of a velocity field, $\|v\|^2 = \int (Lv(x), v(x))dx$. The dual space of the Lie algebra, $\mathfrak{g}^*$ consists of vector-valued distributions. The velocity, $v \in \mathfrak{g}$, maps to its dual deformation momenta, $m \in \mathfrak{g}^*$, via the operator L such that $m = Lv$. Using this norm geodesics are defined as energy minimizing paths between their endpoints. The distance between id and diffeomorphism ϕ is defined via the optimization problem:

$$d(id,\phi)^2 = \inf\left\{\int_0^1 \|v(\cdot,t)\|^2 dt, \text{ subject to: } \varphi(\cdot,1) = \phi\right\}.$$

EPDiff for geodesic evolution: Given the initial velocity, $v_0 \in \mathfrak{g}$, or equivalently, the initial momentum, $m(0) = m_0 \in \mathfrak{g}^*$, the geodesic path $\varphi(t)$ satisfies the EPDiff equation [1, 8]:

$$\frac{d}{dt}m = -\operatorname{ad}_v^* m = -(Dv)^T m - Dmv - (\operatorname{div} v)m \tag{1}$$

where D denotes the Jacobian matrix, and the operator ad^* is the dual of the negative Jacobi-Lie bracket of vector fields [8, 1, 10]: $\mathrm{ad}_v\, w = -[v, w] = (Dv)w - (Dw)v$.

3 Polar Factorization of Diffeomorphisms and PDiff: the Space of Irrotational Diffeomorphisms

Brenier's [3] polar factorization of diffeomorphisms states that any diffeomorphism φ of $\mathbb{R}^d$ can be uniquely written as a composition

$$\varphi = S \circ \psi, \quad \text{where } \psi = \nabla \rho \tag{2}$$

for some convex function $\rho : \mathbb{R}^d \to \mathbb{R}$ and where $S \in \mathrm{SDiff}(\mathbb{R}^d)$ is a measure-preserving diffeomorphism. This decomposition is analogous to the classical polar factorization of matrices. Just as an invertible matrix can be written as product of a positive definite matrix and a unitary matrix, the Jacobian of the deformation ψ, $D\psi = H\rho$ is the Hessian of the convex function ρ, and as such is a symmetric positive-definite matrix, and as S is volume-preserving its Jacobian DS has determinant 1 every where and is unitary. Brenier's polar factorization of $\mathrm{Diff}(\mathbb{R}^d)$ is intimately connected to the Helmholtz-Hodge decomposition of vector fields, which has proven useful for modeling incompressible deformation in computational anatomy [6]. The Helmholtz-Hodge decomposition states that any compactly supported square-integrable C^2 vector field $v \in \mathfrak{g}$ can be written as

$$v = \nabla f + \operatorname{curl} A \tag{3}$$

where $f \in H^1(\mathbb{R}^d)$ and $A \in H^{\mathrm{curl}}(\mathbb{R}^d)$. This constitutes a decomposition of $\mathfrak{g}$ into two linear subspaces: one containing irrotational vector fields represented as gradients of scalar Sobolev functions and one containing incompressible (divergence-free) vector fields represented as curls of Sobolev vector fields. We will denote by $\mathfrak{g}_P$ the subspace of irrotational vector fields and by $\mathfrak{g}_S$ the subspace of divergence-free vector fields, so that $\mathfrak{g} = \mathfrak{g}_P \oplus \mathfrak{g}_S$. Further more assuming compact support, f is uniquely determined by the divergence of v and f satisfies the Poisson equation:

$$\operatorname{div}(v) = g, \quad \Delta f = g, \tag{4}$$

where Δ is the Laplacian operator and $g \in L^2(\mathbb{R}^d)$. Define the space $\mathrm{PDiff}(\mathbb{R}^d)$ to be the space of diffeomorphisms $\psi(1)$ for which there exists a smooth path $\psi(t)$ of diffeomorphisms satisfying

$$\psi(0) = id \quad \text{and} \quad \frac{d}{dt}\psi(t) = v(t) \circ \psi(t) \qquad \forall t \in [0, 1] \tag{5}$$

for some time-varying collection of $v(t) \in \mathfrak{g}_P$. This implies that every $\phi(1) \in \mathrm{PDiff}(\mathbb{R}^d)$ is determined by a time-varying H^1 scalar field $f(t)$:

$$\frac{d}{dt}\psi(t) = (\nabla f(t)) \circ \psi(t). \tag{6}$$

By Liouville's theorem, all such ψ have symmetric positive-definite Jacobian matrices and can be written as the gradient of a convex function. This definition mimics that of $\mathrm{SDiff}(\mathbb{R}^d)$, the space of compactly-supported incompressible diffeomorphisms of $\mathbb{R}^d$. It is to be noted that although $\mathrm{SDiff}(\mathbb{R}^d)$ is a subgroup of diffeomorphisms of $\mathbb{R}^d$, $\mathrm{PDiff}(\mathbb{R}^d)$ is a not a subgroup as $\mathfrak{g}_P$ is not closed with respect to Lie bracket. Just as the with symmetric positive definite matrices, composition of two irrotational diffeomorphisms is not necessarily an irrotational diffeomorphism.

4 Metric and Geodesics on $\mathrm{PDiff}(\mathbb{R}^d)$

The Helmholtz-Hodge decomposition assumes that the divergence of v is square-integrable, so the most natural inner product to induce on irrotational vector fields is the one induced via the L^2 inner product on it divergence, the natural inner product on $\mathfrak{g}_P$ becomes:

$$\langle v, w\rangle_{\mathfrak{g}_P} = \langle \operatorname{div} v, \operatorname{div} w\rangle_{L^2} = \langle \operatorname{div} \nabla f, \operatorname{div} \nabla h\rangle_{L^2} = \int_{\Omega} \Delta f(x) \Delta h(x) dx \tag{7}$$

where Δ is the Laplacian operator. Notice that with the above inner product if g is the divergence of v, the norm of v is simply the L^2 norm of g:

$$\|v\|^2_{\mathfrak{g}_P} = \|\operatorname{div} v\|^2_{L^2(\mathbb{R}^d)} = -\int (\nabla \operatorname{div} v(x))^T v(x) dx = \|g\|^2_{L^2(\mathbb{R}^d)}. \tag{8}$$

This is the $\dot{H}^1$metric, $\langle -\Delta v, w\rangle$, restricted to $\mathfrak{g}_P$, which follows from the fact that $\operatorname{curl} v = 0$ in $\mathfrak{g}_P$ and the identity $\Delta v = \nabla \operatorname{div} v - \operatorname{curl}\operatorname{curl} v$.

Letting $\psi_1, \psi_2 \in \mathrm{PDiff}(\Omega)$ be two diffeomorphisms, a geodesic between ψ_1 and ψ_2 is a path $\alpha(t) \in \mathrm{PDiff}(\mathbb{R}^d)$ connecting ψ_1, ψ_2 that minimizes

$$S(\alpha) = \frac{1}{2}\int_0^1 \|\dot{\alpha}(t)\|^2 dt = \frac{1}{2}\int_0^1 \int |(\Delta f(t))(x)|^2 dx dt. \tag{9}$$

Geodesics on $\mathrm{PDiff}(\mathbb{R}^d)$ are actually minimizing curves in all of $\mathrm{Diff}(\mathbb{R}^d)$ with the constraint that the right-trivialized velocity lie in $\mathfrak{g}_P$ at all times, and are sub-Riemannian geodesics on the Lie group $\mathrm{Diff}(\mathbb{R}^d)$. The theory of sub-Riemannian geodesics in Lie groups has been studied previously [5]. We define the momentum associated with the velocity v as $m = -\nabla \operatorname{div} v = -\nabla g$. Geodesics in $\mathrm{PDiff}(\mathbb{R}^d)$ satisfy the constrained Euler-Poincaré equation 1 with the constraint that v is curl free. Substituting $m = -\nabla g, v = \nabla f$ and $\operatorname{div} v = g$ the constrained Euler-Poincaré equation simply becomes:

$$\frac{d}{dt}\nabla g = -Hg\nabla f - (Hf)^T \nabla g - g\nabla g. \tag{10}$$

The Hessian matrix is always symmetric and notice that $\nabla(g^2) = 2g\nabla g$, so we can rewrite this using the product rule as

$$\nabla \frac{d}{dt} g = -\nabla\left(\nabla g^T \nabla f + \frac{1}{2}g^2\right). \tag{11}$$

Along with our boundary conditions on g this implies that

$$\dot{g} + \nabla g^T v = -\frac{1}{2} g^2. \tag{12}$$

The left-hand side has the form of a material derivative, suggesting a change to Lagrangian coordinates. Introducing $\gamma(t) = g \circ \psi(t)$, implying $\dot{\gamma} = \dot{g} \circ \psi + ((\nabla g)^T v) \circ \psi$ we see that

$$\dot{\gamma}(t) = -\frac{1}{2}\gamma(t)^2, \quad \text{or} \quad \gamma(t) = \frac{\gamma(0)}{\frac{1}{2}t\gamma(0) + 1} \quad . \tag{13}$$

Using the shorthand $g_0 = g(0)$ and the assumption $\psi(0) = id$, we arrive at

$$g(t) \circ \psi(t) = \frac{g_0}{\frac{1}{2}tg_0 + 1}. \tag{14}$$

The quantity $g(t)$ is, by definition, the divergence of the velocity at time t. Using the well-known Liouville's formula we relate this directly to the determinant of the Jacobian matrix of the diffeomorphism ψ as follows:

$$|D\psi(t)| = \exp \int_0^t (\operatorname{div} v(s)) \circ \psi(s) ds = \exp \int_0^t g(s) \circ \psi(s) ds \tag{15}$$

$$= \exp \int_0^t \frac{g_0}{\frac{1}{2}sg_0 + 1} ds = \left(\frac{1}{2}tg_0 + 1\right)^2. \tag{16}$$

Using the solution of the EPDiff equation we can explicitly write the expression for the distance in $\mathrm{PDiff}(\mathbb{R}^d)$ between the identity and any irrotational diffeomorphism ψ. As the metric is simply the L^2 norm of g, by conservation of momenta along a geodesic we have

$$d(id, \psi)^2 = \|g_0\|^2_{L^2(\mathbb{R}^d)} = 4 \int_{\mathbb{R}^d} (\sqrt{|D\psi|} - 1)^2 dx. \tag{17}$$

The simplicity of the above formula comes from the fact that by solving the EPDiff equation, g_0 is essentially the log map on $\mathrm{PDiff}(\mathbb{R}^d)$ with the $\dot{H}^1$metric.

5 Curvature of $\mathrm{PDiff}(\mathbb{R}^d)$

We now use the relationship between g_0 and $|D\psi|$ to show that the curvature of $\mathrm{PDiff}(\mathbb{R}^d)$ with the $\dot{H}^1$metric is 0. Define the following mapping from ψ to the divergence of its initial velocity field:

$$P : \mathrm{PDiff}(\mathbb{R}^d) \to L^2(\mathbb{R}^d) \tag{18}$$

$$P(\psi) = 2(\sqrt{|D\psi|} - 1) = g_0. \tag{19}$$

We first need the following Lemma:

Lemma 1 *The pushforward of a vector field* $u \circ \psi \in T_\psi \mathrm{PDiff}(\mathbb{R}^d)$ *under the mapping* P *is given by the formula*

$$TP(u \circ \psi) = \sqrt{|D\psi|} (\operatorname{div} u) \circ \psi. \tag{20}$$

Proof. Let ψ_s be a family of irrotational diffeomorphisms indexed by the real variable s and satisfying

$$\psi_0 = \psi, \quad \frac{d}{ds}|_{s=0}\psi_s = u \circ \psi. \tag{21}$$

Then the pushforward of the vector field u is defined as

$$TP(u \circ \psi) = \frac{d}{ds}|_{s=0} P\psi_s. \tag{22}$$

A straightforward computation then yields

$$TP(u \circ \psi) = 2\frac{d}{ds}|_{s=0}\sqrt{|D\psi_s|} = \frac{1}{\sqrt{|D\psi|}}\frac{d}{ds}|_{s=0}|D\psi_s| = \sqrt{|D\psi|}(\operatorname{div} u) \circ \psi. \tag{23}$$

Theorem 1 *The mapping P is an isometry from* $\mathrm{PDiff}(\mathbb{R}^d)$ *into an open subset of* $L^2(\mathbb{R}^d)$.

Proof. As the pushfoward is only zero for divergence-free vector fields, Lemma 1 directly implies that P is injective on $\mathrm{PDiff}(\mathbb{R}^d)$. To prove that P is furthermore an isometry, we compute the pullback of the L^2 metric for any two vector fields $u \circ \psi, w \circ \psi \in T_\psi \mathrm{PDiff}(\mathbb{R}^d)$:

$$\langle u, w\rangle_{P^*} = \langle TP(u \circ \psi), TP(w \circ \psi)\rangle_{L^2(\mathbb{R}^d)}. \tag{24}$$

Plugging in and performing a change of variables, we have

$$\begin{aligned}
\langle u, w\rangle_{P^*} &= \int \sqrt{|D\psi(x)|}(\operatorname{div} u) \circ \psi(x)\sqrt{|D\psi(x)|}(\operatorname{div} w) \circ \psi(x)dx && (25)\\
&= \int |D\psi(x)|(\operatorname{div} u) \circ \psi(x)(\operatorname{div} w) \circ \psi(x)dx && (26)\\
&= \int \operatorname{div} u(y) \operatorname{div} w(y)dy && (27)\\
&= \langle \operatorname{div} u, \operatorname{div} v\rangle_{L^2(\mathbb{R}^d)}, && (28)
\end{aligned}$$

which is our right-invariant metric on $\mathrm{PDiff}(\mathbb{R}^d)$, proving that P is a local isometry. By the uniqueness of Brenier's polar factorization, the mapping P is injective, completing the proof. □

The property that P is an isometry is remarkable in that it implies (since $L^2(\mathbb{R}^d)$ is a flat vector space) that with the $\dot{H}^1$metric, $\mathrm{PDiff}(\mathbb{R}^d)$ has zero Riemannian curvature[1]. Another important consequence is that under P, geodesics in $\mathrm{PDiff}(\mathbb{R}^d)$ map to straight lines in $L^2(\mathbb{R}^d)$. The image of P consists of all L^2

[1] This has been observed very recently in [2] for the special case of $d = 1$ where $\mathrm{PDiff}(\mathbb{R}^1) = \mathrm{Diff}(\mathbb{R}^1)$ as the only compactly-supported measure-preserving diffeomorphism of the real line is the identity mapping.

functions with values strictly greater than -2, implying that geodesics can leave this open subset in finite time. Given an initial velocity field, this blow-up time is determined by the minimum value of its divergence g_0 and Eq. 16.

The P map is injective, so given $g_0 \in L^2(\mathbb{R}^d)$, there is a unique irrotational diffeomorphism $\psi \in \mathrm{PDiff}(\mathbb{R}^d)$ in the inverse image $P^{-1}(g_0)$. Computation of ψ is equivalent to computing the exponential map in $\mathrm{PDiff}(\mathbb{R}^d)$. We are unaware of a closed-form method for computing ψ, but it may be computed numerically using Eq. 14 to compute $g(t) = \mathrm{div}\, v(t)$ at each time, then solving for the velocity field $v(t)$ and integrating the flow.

6 Irrotational Image Registration

Consider a registration problem in which two images $I_0, I_1 \in L^2(\mathbb{R}^d)$ are given and one wishes to find an irrotational deformation $\psi \in \mathrm{PDiff}(\mathbb{R}^d)$ that best matches the two images. Analogous to the LDDMM approach, we introduce the energy functional

$$E(\psi) = \frac{1}{2\sigma^2}\|I_0 \circ \psi^{-1} - I_1\|^2_{L^2(\mathbb{R}^d)} + d(id, \psi)^2 \tag{29}$$

where d denotes the geodesic distance within $\mathrm{PDiff}(\mathbb{R}^d)$. However, unlike with general LDDMM, the distance term can now be evaluated in closed form only using ψ:

$$E(\psi) = \frac{1}{2\sigma^2}\|I_0 \circ \psi^{-1} - I_1\|^2_{L^2(\mathbb{R}^d)} + 4\|\sqrt{|D\psi|} - 1\|^2_{L^2(\mathbb{R}^d)}. \tag{30}$$

This allows us to take the Sobolev variation of E with respect to ψ directly by first taking the L^2 variation and then sharping it using the inverse of the metric. Let $\nabla c \in \mathfrak{g}_P$ be a perturbation of ψ, and let $\psi_s \in \mathrm{PDiff}(\mathbb{R}^d)$ be a family of irrotational diffeomorphisms parametrized by the real variable s, satisfying

$$\psi_0 = \psi \quad \text{and} \quad \frac{d}{ds}|_{s=0}\psi_s = (\nabla c) \circ \psi. \tag{31}$$

Then the variation of E with respect to ψ in the direction of ∇c is computed via

$$(\delta E, \nabla c) = \frac{d}{ds}|_{s=0} E(\psi_s) \tag{32}$$

$$= \frac{d}{ds}|_{s=0} \frac{1}{2\sigma^2}\|I_0 \circ \psi_s^{-1} - I_1\|^2_{L^2} + 4\|(\sqrt{|D\psi_s|} - 1)\|^2_{L^2} \tag{33}$$

$$= \frac{d}{ds}|_{s=0} \frac{1}{2\sigma^2}\int_\Omega (I_0 \circ \psi_s^{-1}(y) - I_1(y))^2 dy + 4\int_\Omega (\sqrt{|D\psi_s(x)|} - 1)^2 dx \tag{34}$$

$$= \frac{1}{\sigma^2}\int_\Omega (I_0 \circ \psi^{-1}(y) - I_1(y))\nabla(I_0 \circ \psi^{-1}(y))^T \nabla c(y) dy \tag{35}$$

$$+ 4\int_\Omega (\sqrt{|D\psi(x)|} - 1)\frac{1}{\sqrt{|D\psi(x)|}}\frac{d}{ds}|_{s=0}|D\psi_s(x)| dx. \tag{36}$$

Using $\frac{d}{ds}|_{s=0}|D\psi_s(x)| = (\operatorname{div} \nabla c) \circ \psi(x)|D\psi(x)|$ and the fact that, for compactly supported vector fields, the adjoint of the divergence is the negative gradient, we have

$$(\delta E, \nabla c) = -\frac{1}{\sigma^2} \int_\Omega \operatorname{div} \left((I_0 \circ \psi^{-1}(y) - I_1(y)) \nabla (I \circ \psi^{-1}(y)) \right) c(y) dy \tag{37}$$

$$+ 4 \int_\Omega (\sqrt{|D\psi| \circ \psi^{-1}(y)} - 1) \sqrt{|D\psi| \circ \psi^{-1}(y)} \Delta c(y) |D\psi^{-1}(y)| dy. \tag{38}$$

Now we use the identity $(D\psi^{-1}) \circ \psi(x) = (D\psi)^{-1}(x)$ and self-adjointness of the Laplacian to simplify this to

$$(\delta E, \nabla c) = -\frac{1}{\sigma^2} \int_\Omega \operatorname{div} \left((I_0 \circ \psi^{-1}(y) - I_1(y)) \nabla (I_0 \circ \psi^{-1}(y)) \right) c(y) dy \tag{39}$$

$$+ 4 \int_\Omega c(y) \Delta \left(1 - \sqrt{|D\psi^{-1}(y)|} \right) dy. \tag{40}$$

By adjointing the gradient in the left-hand side we see that since this must hold for all c, we have

$$\operatorname{div} \delta E = \frac{1}{\sigma^2} \operatorname{div} \left((I_0 \circ \psi^{-1} - I_1) \nabla (I_0 \circ \psi^{-1}) \right) + 4 \Delta (\sqrt{|D\psi^{-1}|} - 1). \tag{41}$$

In order to convert δE to the Sobolev variation of E, we solve the following for the scalar function b:

$$\Delta^2 b = \frac{1}{\sigma^2} \operatorname{div} \left((I_0 \circ \psi^{-1} - I_1) \nabla (I_0 \circ \psi^{-1}) \right) + 4 \Delta (\sqrt{|D\psi^{-1}|} - 1) \tag{42}$$

then update ψ via

$$\psi(x) \mapsto \psi(x) - \epsilon (\nabla b) \circ \psi(x) \tag{43}$$

for some step-size ϵ. In practice, as ψ is never needed we directly update only ψ^{-1} via

$$\psi^{-1}(y) \mapsto \psi^{-1}(y + \epsilon \nabla b(y)). \tag{44}$$

Notice that this allows ψ^{-1} to be optimized directly in a gradient-based scheme without the need for numeric integration of geodesic equations or adjoint equations.

7 Symmetric Image Registration

In this section, we present an image registration approach that is symmetric with respect to swapping of the input images. Consider re-weighting the image match term by the square root of the Jacobian determinant.

$$E(\psi) = \frac{1}{2\sigma^2} \int |I_0 \circ \psi^{-1}(y) - I_1(y)|^2 \sqrt{|D\psi^{-1}(y)|} dy + d(id, \psi)^2. \tag{45}$$

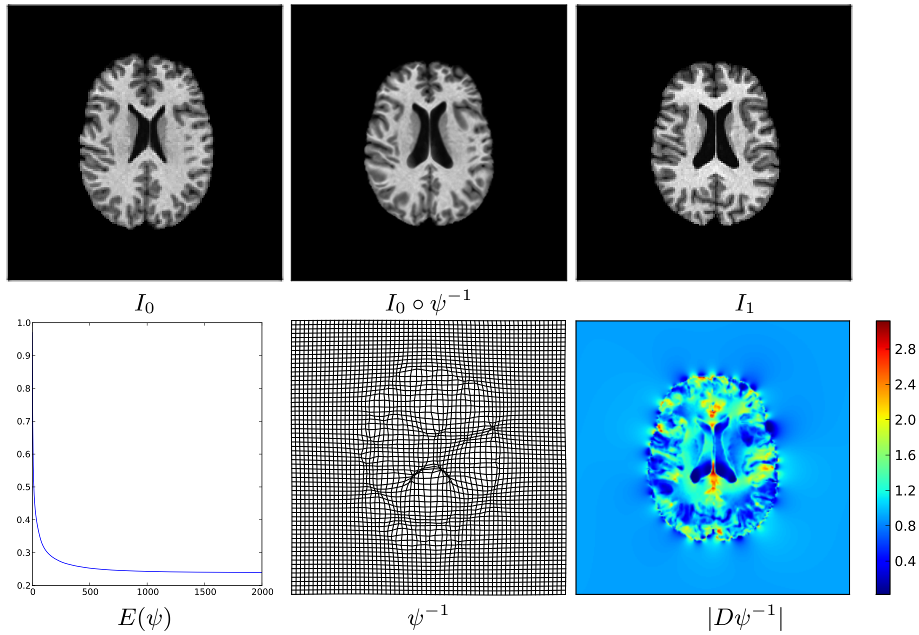

Fig. 1. Neuroimaging study, symmetric irrotational registration results. The algorithm was run with inputs I_0, I_1 to generate the irrotational diffeomorphism ψ. The plot of energy $E(\psi)$ at each iteration is shown on the left in the lower column showing good convergence, along with the estimated deformation ψ and its Jacobian determinant.

Now using the change of variables $x = \psi^{-1}(y)$

$$E(\psi) = \frac{1}{2\sigma^2} \int \|I_0(x) - I_1 \circ \psi(x)\|^2 \sqrt{|(D\psi^{-1}) \circ \psi(x)|} |D\psi(x)| dx + d(\psi^{-1}, id)^2. \tag{46}$$

Using the inversion-invariance of our metric we rewrite the cost functional as

$$E(\psi) = \frac{1}{2\sigma^2} \int |I_0(x) - I_1 \circ \psi(x)|^2 \sqrt{|D\psi(x)|} dx + d(id, \psi)^2. \tag{47}$$

This has the same form as the original function in which the first image I_0 was deformed to match I_1, but instead we match I_1 to I_0. So the introduction of the square-root Jacobian determinant into the image match term has the effect of making the image registration problem invariant under relabeling of the input images. Computing the variation of this functional is very similar to the method in the previous section, and leads us to the following biharmonic equation:

$$\Delta^2 b = \frac{1}{\sigma^2} \operatorname{div} \left((I_0 \circ \psi^{-1} - I_1) \sqrt{|D\psi^{-1}|} \nabla (I_0 \circ \psi^{-1}) \right) \tag{48}$$

$$+ \Delta \left(\left(4(1 - \sqrt{|D\psi^{-1}|}) - \frac{1}{4\sigma^2} |I_0 \circ \psi^{-1} - I_1|^2 \right) \sqrt{|D\psi^{-1}|} \right). \tag{49}$$

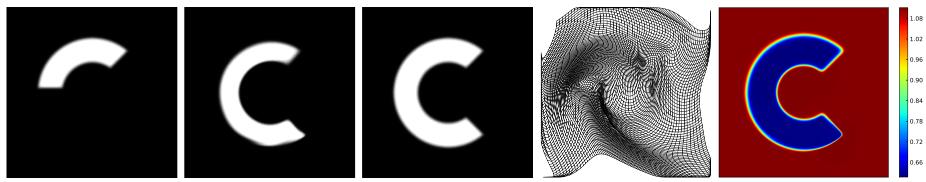

I_0 $\quad I_0 \circ \varphi^{-1}$ $\quad I_1$ $\quad \varphi^{-1}$ $\quad |D\varphi^{-1}|$

Fig. 2. Synthetic study, symmetric hybrid image registration results. Shown here are the input image I_0, along with the deformed image $I_0 \circ \varphi^{-1}$, the target image I_1, the deformation φ^{-1}, and its Jacobian determinant $|D\varphi^{-1}|$.

After solving this equation for b, we take the gradient then update ψ just as we did in the asymmetric case.

Neuroimaging Study We have implemented the symmetric irrotational image registration algorithm and applied it to two structural MRI images. Figure 1 shows the result of symmetric irrotational image registration. Notice that even without allowing any local rotation, the two images are matched quite well. In the bottom row is shown the energy at each iteration, indicating very stable convergence. Also notice that the Jacobian determinant clearly indicating regions of expansion and contraction. In our irrotational matching method, the Jacobian determinant entirely characterizes the diffeomorphism.

8 Pseudo-Riemannian Hybrid Irrotational/Incompressible Registration

In this section we present an extension of the irrotational-only algorithms which allows an incompressible component to be estimated without any penalty. The right invariant metric on $\mathrm{PDiff}(\mathbb{R}^d)$ defined in Section 4 is also a right-invariant Riemannian pseudo-metric on all of $\mathrm{Diff}(\mathbb{R}^d)$. The null space of this metric consists of precisely all divergence-free vector fields, which is the Lie algebra $\mathfrak{g}_S$ associated with the group of volume-preserving diffeomorphisms $\mathrm{SDiff}(\mathbb{R}^d)$. Consider registration using a general diffeomorphism $\varphi = S \circ \psi$, where $S \in \mathrm{SDiff}(\mathbb{R}^d)$. Using the polar factorization of $\mathrm{Diff}(\mathbb{R}^d)$, we replace $E(\psi)$ with the functional

$$E(\psi, S) = \frac{1}{2\sigma^2} \int |I_0 \circ \psi^{-1} \circ S^{-1}(y) - I_1(y)|^2 \sqrt{|D\psi^{-1} \circ S^{-1}(y)|} dy + d(id, \psi)^2. \tag{50}$$

Equation 50 is rewritten in terms of φ^{-1} only, using the fact that $|DS| = 1$ everywhere:

$$E(\varphi) = \frac{1}{2\sigma^2} \int |I_0 \circ \varphi^{-1}(y) - I_1(y)|^2 \sqrt{|D\varphi^{-1}(y)|} dy + 4 \int (\sqrt{|D\varphi^{-1}(y)|} - 1)^2 dy. \tag{51}$$

This is optimized by decomposing the Sobolev variation of $E(\varphi)$ using the Helmholtz-Hodge decomposition into irrotational and incompressible components, then performing gradient descent steps in either component. The irrotational updates are performed exactly as described in the previous section, and since the incompressible updates do not effect the Jacobian determinant, the incompressible update direction $w \in \mathfrak{g}_S$ is found by simply solving

$$\Delta w = -\frac{1}{\sigma^2}\left(I_0 \circ \varphi^{-1} - I_1\right)\sqrt{|D\varphi^{-1}|}\nabla(I_0 \circ \varphi^{-1}) \tag{52}$$

and projecting onto the space of divergence-free vector fields $\mathfrak{g}_S$. This projection has been discussed previously in the literature and is performed efficiently in the Fourier domain while simultaneously solving the above Poisson's equation [6].

Synthetic example In order to test the performance of our algorithms in the presence of large deformation, a simulated experiment was also performed. Two synthetic two-dimensional datasets were generated, simulating a completed "C" and a half C. In Fig. 2 are shown the results of the hybrid image registration algorithm. Notice that the deformed half C image, $I_0 \circ \varphi^{-1}$, agrees very well with the full C image, I_1 and that this is achieved while maintaining a diffeomorphic transformation. Finally, as we penalize the L^2 norm of the square root Jacobian determinant, the Jacobian determinant of the overall deformation is distributed very evenly across the entire deforming region, instead of being concentrated at a single advancing edge.

9 Multi-Scale and Scale Independence of the Metric.

Along the way, we have noted that the Helmholtz-Hodge decomposition leads us naturally to the $\dot{H}^1$metric we use, which is essentially the Laplacian metric restricted to $\mathrm{PDiff}(\mathbb{R}^d)$. This metric is also quite natural in the sense that it is spatial scale-independent. The Green's function of the Laplacian is given by

$$K(x,y) = \frac{1}{|x-y|}.$$

Clearly then, any scaling of the domain has the simple effect of multiplying this kernel by a constant factor, or equivalently changing the speed of the geodesics. By contrast, the commonly used Sobolev metric $s-\Delta$ has as its Green's function

$$K(x,y) = e^{-s|x-y|}. \tag{53}$$

Changing of the spatial scale in this case has the effect of changing the bandwidth and in fact changes the curvature of the space. The Laplacian metric allows phenomena at all scales to influence the image registration. This follows from the recently-developed theory of multi-scale image registration [4, 9] along with the fact that the Laplacian kernel can be written as an integral of $s - \Delta$ kernels as follows:

$$\frac{1}{|x-y|} = \int_0^\infty e^{-s|x-y|} ds. \tag{54}$$

10 Discussion

We have shown that Brenier's polar decomposition of compactly-supported diffeomorphisms, along with the divergence metric on the irrotational component leads to novel new image registration algorithms. Furthermore, Theorem 1 shows that with this metric, the $\mathrm{PDiff}(\mathbb{R}^d)$ component can be isometrically embedded in the flat vector space $L^2(\mathbb{R}^d)$, a fact that underlies the efficiency of our new algorithms. Even more importantly, it has far reach statistical implications, allowing statistics to be performed in $\mathrm{PDiff}(\mathbb{R}^d)$ without the difficulties that often accompany statistics on curved manifolds. In particular, parallel transport of a vector field $w \in \mathfrak{g}_P$ along a curve in $\psi(t) \in \mathrm{PDiff}(\mathbb{R}^d)$ is path-independent and can be conveniently computed in closed form using only the divergence $h(t) = \operatorname{div} w(t)$ and the diffeomorphism at times 0 and 1:[2]

$$h(1) = \sqrt{\frac{|D\psi^{-1}(1)|}{|D\psi^{-1}(0)|}} h(0). \tag{55}$$

Flatness also enables simplification of other intrinsic methods involving the covariant derivative and curvature tensor such as geodesic regression and Jacobi fields, principal geodesic analysis, as well as Riemannian polynomials and splines.

Acknowledgements

The authors would like to thank Xavier Pennec and Marco Lorenzi for discussions about irrotational diffeomorphisms, as well as Peter Michor and Martin Bauer for invaluable discussions regarding the flatness of diffeomorphism spaces on $\mathbb{R}^1$, all of which occurred primarily at the Workshop on Geometry and Statistics 2012 organized by University of Copenhagen and Aarhus University and was the inspiration for this work. This work was in part supported by NIH grants 5R01EB007688, P41 RR023953 and 5R21HL110059-02.

References

1. Arnol'd, V.I.: Sur la géométrie différentielle des groupes de Lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits. Ann. Inst. Fourier 16, 319–361 (1966)
2. Bauer, M., Bruveris, M., Michor, P.W.: The homogeneous Sobolev metric of order one on diffeomorphism groups on the real line. arXiv preprint math-ap/12092836 (Sep 2012)
3. Brenier, Y.: Polar factorization and monotone rearrangement of vector-valued functions. Communications on pure and applied mathematics 44(4), 375–417 (1991)
4. Bruveris, M., Risser, L., Vialard, F.X.: Mixture of kernels and iterated semi-direct product of diffeomorphism groups. arXiv preprint arXiv:1108.2472 (2011)
5. Fedorov, Y., Jovanović, B.: Integrable nonholonomic geodesic flows on compact Lie groups. arXiv preprint math-ph/0408037 pp. 115–152 (2006)
6. Hinkle, J., Szegedi, M., Wang, B., Salter, B., Joshi, S.: 4D CT image reconstruction with diffeomorphic motion model. Medical Image Analysis 16(6), 1307–1316 (Aug 2012)

[2] Due to space limitations, we omit the derivation of this formula, which follows from Theorem 1 and Lemma 1.

7. Lorenzi, M., Ayache, N., Pennec, X.: Regional flux analysis of longitudinal atrophy in Alzheimers disease. Medical Image Computing and Computer-Assisted Intervention–MICCAI 2012 pp. 739–746 (2012)
8. Miller, M.I., Trouvé, A., Younes, L.: Geodesic shooting for computational anatomy. Journal of Mathematical Imaging and Vision 24(2), 209–228 (2006)
9. Sommer, S., Lauze, F., Nielsen, M., Pennec, X.: Kernel bundle EPDiff: evolution equations for multi-scale diffeomorphic image registration. Scale Space and Variational Methods in Computer Vision pp. 677–688 (2012)
10. Younes, L., Arrate, F., Miller, M.I.: Evolutions equations in computational anatomy. NeuroImage 45(1), S40–S50 (2009)