Surface Matching via Currents

Marc Vaillant and Joan Glaunès

[1] CIS, Johns Hopkins University, Baltimore, MD. marc@jhu.edu
[2] LAGA, Université Paris 13, Villetaneuse, FRANCE.
glaunes@math.univ-paris13.fr

Abstract. We present a new method for computing an optimal deformation between two arbitrary surfaces embedded in Euclidean 3-dimensional space. Our main contribution is in building a norm on the space of surfaces via representation by currents of geometric measure theory. Currents are an appropriate choice for representations because they inherit natural transformation properties from differential forms. We impose a Hilbert space structure on currents, whose norm gives a convenient and practical way to define a matching functional. Using this Hilbert space norm, we also derive and implement a surface matching algorithm under the large deformation framework, guaranteeing that the optimal solution is a one-to-one regular map of the entire ambient space. We detail an implementation of this algorithm for triangular meshes and present results on 3D face and medical image data.

1 Introduction

Surfaces embedded in 3D are important geometric models for many objects of interest in image analysis. They are often the appropriate abstractions for studying gross shape, either because the structure of interest is inherently 2D (e.g. the outer cortex of the human brain, human face, etc.), or because its shape can be efficiently and completely captured by its bounding surface (e.g. planes, animals, or anatomical structures of the human body, etc.). A fundamental task in image analysis applications is to perform a non-rigid matching (deformation) between two occurrences of the same structure. For example, it has been recognized as early as 1917 by D'Arcy Thompson [1], that given representations of a particular anatomic structure in two subjects, an appropriate methodology for comparing their gross morphological differences is to study a *transformation*–uniquely characterized by a natural optimality property–from one structure into the other.

Surface matching is usually achieved via semi-automated procedures in which a small number of substructures, such as landmark points or curves, are identified by hand and then used to guide the transformation of the entire surface [2–6]. One interpretation of the problem is to consider surface matching as a "point correspondence" task: for each point on the discretized template surface find its corresponding point on the target. Fully automated approaches to this problem have been developed in [7]. However, a fundamental issue with this point of view is that, due to discretization, a point on one surface need not have a homologous point on the other. This problem is handled in [7] by simultaneously identifying and rejecting points with no corresponding pair as

"outliers". A second issue is that geometric information is necessarily discarded when reducing surfaces–inherently 2D objects–to 0-dimensional point sets. Another related approach is the work of Wang et. al. [8]. This approach does use local geometric constraints by including surface curvature in their matching criterion. Its advantage is that both triangulation and correspondence is established simultaneously. Another elegant approach includes the work of Davies et. al. [9] in which the correspondence problem is tackled by building the "best" model via a set of optimality criteria.

We develop a surface matching approach in which the two issues mentioned above are overcome naturally in the fundamental theoretical framework. Our approach follows most closely the work of [10], but differs in that we represent surfaces as the generalized distributions of deRham called **currents** [11], instead of the classical distributions of Schwartz. As in [10], distribution representations allow us to get away from a strict pointwise representation of surfaces and therefore enable us to treat the problem as true surface matching without the point correspondence issue. Furthermore, representation via currents captures the geometry of the structure because it is sensitive to both location and to the first order local geometric structure. We are therefore able to overcome the two issues above via the natural choice of current representation. In this paper we provide a detailed theoretical development of a norm on surfaces which can be used as a matching criterion for finding an optimal transformation from one surface into another. We present one such variational optimization problem under the large deformation diffeomorphism framework. Finally, we derive a discretized version of this variation problem, detail its implementation, and provide results on 3D face and medical image data.

2 Surfaces as currents

In order to build a surface matching algorithm, we need a criterion that measures how "close" or similar one surface is to another. Our strategy is to represent surfaces as objects in a linear space and then to equip this space with a computable norm. The generalized distributions from geometric measure theory, called currents, will serve as the representers. In this section we set the notation and introduce currents as representations of surfaces. The main power, and hence motivation, for using these representations will become clear in Section 2.2 where we see that they are preserved under coordinate transformations.

2.1 2-forms and Currents

The paradigm of the approach is built from mathematical objects called differential m-forms. Although the theory is more general than presented here, we restrict the discussion to the setting of interest: surfaces embedded in $\mathbb{R}^3$. In this setting, we need only introduce differential 2-forms. A differential **2-form** on $\mathbb{R}^3$ is a differential mapping $x \mapsto \omega(x)$ such that for each $x \in \mathbb{R}^3$, $\omega(x)$ is a skew-symmetric bilinear function on $\mathbb{R}^3$. A 2-form is a natural object to be integrated over an oriented smooth surface S because, as we will see, it automatically transforms correctly under a change of coordinates. For each $x \in S$, let u_x^1, u_x^2 be an orthonormal basis of the tangent plane at x.

Abusing notation slightly, we then associate to S the function

$$S(\omega) = \int_S \omega(x)(u_x^1, u_x^2) d\sigma(x), \tag{1}$$

where $d\sigma$ is the element of surface area. Thus, the surface S is seen as a linear functional on the space of 2-forms via (1). More generally, the space of 2-dimensional **currents** is defined as the dual space to C^∞ 2-forms with compact support. It is equipped with the correct topology as in the classical theory of distributions of Schwartz. This definition extends to more singular geometric objects, such as triangular meshes, by replacing the surface measure with 2-dimensional Hausdorff measure. In fact a wide class of geometric subsets of $\mathbb{R}^3$, called rectifiable sets, can be viewed as currents [12].

In the sequel, we continue to abuse notation by using the same letter to denote both a surface as well as its associated representation as a current.

2.2 Push forward of a current

The fundamental property ultimately motivating the representation of surfaces by currents is that it is possible to define an action of diffeomorphisms $\phi : \mathbb{R}^3 \to \mathbb{R}^3$ on currents which coincides with the natural action of ϕ on surfaces (i.e. $S \mapsto \phi(S)$).

First define the **pull back** of a 2-form ω by: $\phi^\sharp \omega(x)(\eta, \nu) = \omega(\phi(x))\left((d_x\phi)\eta, (d_x\phi)\nu\right)$. The **push forward** $\phi_\sharp S$ of a current S is $\phi_\sharp S(\omega) = S(\phi^\sharp \omega)$. The change of coordinates for integration of differential forms [13] states

$$S(\phi^\sharp \omega) = \phi(S)(\omega). \tag{2}$$

That is, $\phi_\sharp S$ is indeed the current associated with $\phi(S)$, which is exactly the natural property we would like our representations to have.

2.3 Vectorial representation

It will be convenient to use a vectorial representation of skew-symmetric bilinear functions on $\mathbb{R}^3$. If B is such a function, its representer $\overline{B} \in \mathbb{R}^3$ satisfies $B(\eta, \nu) = \overline{B} \cdot (\eta \times \nu)$, where $\cdot$ and $\times$ are the euclidean dot and cross products respectively. Therefore a 2-form, $\omega(x)$, will be represented by the vector field $\overline{\omega}(x)$ via this association. Formally, the association between 2-forms and vectors is given by the hodge star operator and duality (see [13]).

2.4 Hilbert space of currents

Recall that the motivation for introducing representations is as a vehicle for constructing a norm on the space of hypersurfaces of $\mathbb{R}^3$. In practice, this norm must be computable. We see in this section that the currents of interest, i.e. those associated with hypersurfaces of $\mathbb{R}^3$ via (1), can be equipped with a Hilbert space structure having an easily computable norm. The machinery of Reproducing kernel Hilbert space (r.k.h.s.) theory is fundamental in this construction (see [14]). Background in the somewhat uncommon setting of differential forms is given next, together with the derivation of the norm.

Let $(W, \langle \cdot, \cdot \rangle_W)$ be a Hibert space of differential 2-forms. The dual space (space of continuous linear functionals) of W is denoted W^*. By the Riesz-Frechet theorem, each $S \in W^*$ has a representer $K_W S \in W$ such that for every $\omega \in W$, $S(\omega) = \langle K_W S, \omega \rangle_W$. K_W is in fact an isometry between W^* and W. We say that W is a r.k.h.s. if for every $x, \xi \in \mathbb{R}^3$, the associated linear **evaluation functional**, δ_x^ξ, defined by $\delta_x^\xi(\omega) = \overline{\omega}(x) \cdot \xi$ belongs to W^*. If W^* is a r.k.h.s., we define the **reproducing kernel** operator k_W by

$$k_W(x, y)\xi = \overline{K_W \delta_x^\xi}(y).$$

Thus it is in fact the reproducing kernel of $\overline{W}$, the space of vector fields corresponding to W. From this definition follows the formula:

$$\langle \delta_x^\xi, \delta_y^\eta \rangle_{W^*} = k_W(x, y)\xi \cdot \eta \tag{3}$$

We impose a slightly stronger constraint than continuity of the evaluation functionals: W is constructed so that it is continuously embedded in the space of continuous bounded 2-forms. That is, there exists some constant c such that $|\omega|_\infty \leq c|\omega|_W$ for all $\omega \in W$. This immediately implies continuity of the evaluation functionals, and furthermore, if S is a surface, we have

$$|S(\omega)| \leq \int_S |\delta_x^{u_x^1 \times u_x^2}(\omega(x))| d\sigma(x) \leq \sigma(S)\, c|\omega|_W.$$

Hence $S \in W^*$, and we are now able to compare submanifolds via the dual space norm on W^*.

3 Surface Matching

Equipped with an appropriate representative space W^*, as described in the previous section, we can now state an optimization problem for mapping one surface into another. We have chosen the well established "large deformation" setting which provides a solution that is a diffeomorphism of the ambient space. This framework is founded in the paradigm of Grenander's group action approach for modeling objects. Abstractly, an appropriate group of transformations, $\mathcal{G}$, is defined together with a group action, which act on a set of objects or structures of interest, $\mathcal{M}$. The idea is to study two elements S_1 and S_2 of $\mathcal{M}$ through an "optimal" transformation $\phi \in \mathcal{G}$ that registers these objects (i.e. $\phi S_1 = S_2$). This approach shifts the focus of the modeling effort onto the study of transformations, as envisioned by D'Arcy Thompson.

In the large deformation setting, $\mathcal{G}$ is a subgroup of diffeomorphisms and the structures of interest in this paper, $\mathcal{M}$, are hypersurfaces of $\mathbb{R}^3$. Optimality is realized by considering all curves $\phi_t \in \mathcal{G}, t \in [0, 1]$ connecting two elements in $\mathcal{G}$ via the group action. The optimal transformation is given by ϕ_1, for the curve which minimizes the accumulated infinitesimal variations in $\mathcal{G}$ through a riemannian structure. We next detail the construction of the group $\mathcal{G}$ and define formally the optimization problem.

3.1 Large deformation framework

The fundamental object of construction is a Hilbert space V, with inner product $\langle \cdot, \cdot \rangle_V$, of smooth vector fields (at least C^1) defined on the background space $\mathbb{R}^3$. For all time dependent families of elements of V, written $v_t \in V$ for $t \in [0,1]$, such that $\int_0^1 |v_t|_V dt < \infty$, the solution ϕ_t at time $t = 1$, of

$$\frac{\partial \phi}{\partial t} = v_t \circ \phi_t, \tag{4}$$

with $\phi_0(x) = x$, is a unique diffeomorphism (see [15, 16]). The collection of all such solutions defines our subgroup of diffeomorphisms $\mathcal{G}_V$, and the inner product $\langle \cdot, \cdot \rangle_V$ equips it with a Riemannian structure. We will sometimes denote ϕ^v for an element of $\mathcal{G}$, explicitly characterizing it by its associated vector field v. The geodesics of $\mathcal{G}$ provide the transformations which match objects in the orbit, and are characterized by extremals of the kinetic energy $\frac{1}{2}\int_0^1 |v_t|_V^2 dt$. In fact, $\mathcal{G}$ can be equipped with a natural right-invariant geodesic distance

$$d_V(\phi, \phi') = \inf \left\{ \Big(\int_0^1 |v_t|_V^2 dt \Big)^{1/2} \,,\ \phi_1^v \circ \phi = \phi' \right\}.$$

3.2 Variational formulation

We define the optimal matching, ϕ_* between two currents S and T as a minimizer of $J_{S,T}(\phi) \doteq d_V(Id, \phi)^2 + |\phi_\sharp S - T|^2_{W^*}/\sigma_R^2$, where σ_R^2 is a trade-off parameter. Equivalently we have $\phi_* = \phi_1^{v_*}$ where v_* is a minimizer of

$$J_{S,T}(v) = \int_0^1 |v_t|_V^2 dt + \frac{1}{\sigma_R^2} |(\phi_1^v)_\sharp S - T|^2_{W^*} \tag{5}$$

The first term of this energy is referred to as the regularizing term, and the second is referred to as the matching, or data attachment term.

In practice, a surface is approximated by a triangular mesh, which also has a current representation. Our strategy is to approximate triangle mesh associated currents in order to derive a computable gradient of the energy (5). We next detail the approximation and gradient derivation.

Let S be a triangular mesh in $\mathbb{R}^3$. Given a face f of S, let f^1, f^2, f^3 denote its vertices, $e^1 = f^2 - f^3, e^2 = f^3 - f^1, e^3 = f^1 - f^2$ its edges, $c(f) = \frac{1}{3}(f^1 + f^2 + f^3)$ its center, and $N(f) = \frac{1}{2}(e^2 \times e^3)$ its normal vector with length equal to its area. We will also denote by S_t the triangular mesh at time t, with faces f_t having vertices $f_t^i = \phi_t(f^i), \quad i = 1, 2, 3$.
The mesh S is represented as a current in the following way

$$S(\omega) = \sum_f \int_f \overline{\omega}(x) \cdot (u_x^1 \times u_x^2) d\sigma_f(x),$$

where σ_f is the surface measure on f. Now, we approximate ω over a face by its value at the center. Thus, we have the approximation $S(\omega) \approx \sum_f \overline{\omega}(c(f)) \cdot N(f)$, so in fact,

the approximation is a sum of linear evaluation functionals $\mathcal{C}(S) = \sum_f \delta^{N(f)}_{c(f)}$, and the matching error can be easily computed using the reproducing kernel as in (3).

From the identity $\phi_\sharp S = \phi(S)$ we can infer two possible approximations to $\phi(S)$:

1. compute the approximation $\mathcal{C}(S)$ and then apply the push forward formula $(\phi_1)_\sharp \mathcal{C}(S)$:

$$\phi_\sharp \delta^\xi_x = \delta^{\det(d_x\phi)(d_x\phi^*)^{-1}\xi}_{\phi(x)} \tag{6}$$

2. first compute S_1 and then compute the approximation $\mathcal{C}(S_1)$.

We have implemented the second approximation. The advantage is that it does not involve the derivatives of ϕ_1, which simplifies the computation of the gradient (cf 3.3). Note, however, that in this case an additional approximation is made since $S_1 \neq \phi_1(S)$. Given either approximation, we can compute explicitly the metric between two surfaces S and T. Let f, g index the faces of S and q, r index the faces of T, the metric $\mathcal{E} = |\mathcal{C}(S) - \mathcal{C}(T)|^2_{W^*}$ between these two surfaces under the second approximation becomes

$$\begin{aligned}\mathcal{E} = &\sum_{f,g} N(f)^t k_W(c(g), c(f))N(g) - 2\sum_{f,q} N(f)^t k_W(c(q), c(f))N(q) \\ &+ \sum_{q,r} N(q)^t k_W(c(q), c(r))N(r).\end{aligned}$$

After a considerable amount of theoretical work, we have arrived at a fairly simple formula which we can analyze intuitively. The first and last terms enforce structural integrity of the two surfaces, while the middle term penalizes geometric and spatial mismatch. Using this approximation we now turn to the computation of the gradient with respect to v_t.

3.3 Gradient of J in $L^2([0,1], V)$

Let x^j index the vertices of S. Like all point-based matching problems in the large deformation setting, it can be shown that the optimal vector fields v_t are of the form

$$v_t(x) = \sum_j k_V(x_j, x)\alpha^j_t, \tag{7}$$

where k_V denotes the reproducing kernel *of the deformation space* V (see [4, 5]). The vectors α^j_t are referred to as **momentum vectors** do to the connection of the large deformation setting to Hamiltonian mechanics (see [17, 18]). It follows from the flow equation that the matching functional (5) is a function only of the trajectories x^j_t.

Gradient of the data attachment term The gradient of the data attachment term, $\mathcal{E}$, in the space $L^2([0,1], V)$ of vector fields is of the form $\nabla\mathcal{E}_t(x) = \sum_j k_V(x^j_t, x) d_{x^j_t}\phi^*_{t1} \nabla_{x^j_t}\mathcal{E}$.

Indeed, for a variation $v_{t,\epsilon} = v_t + \epsilon\widetilde{v}_t$ of the vector field v_t, the corresponding variation of $x_1^j = \phi_1(x^j)$ is (see [19])

$$\widetilde{x}_1^j = \partial_\epsilon x_1^j|_{\epsilon=0} = \int_0^1 d_{x_t^j}\phi_{t1}\widetilde{v}_t(x_t^j)dt,$$

and thus the variation of $\mathcal{E}$ is

$$\begin{aligned}\partial_\epsilon \mathcal{E}|_{\epsilon=0} &= \sum_j \partial_{x_t^j}\mathcal{E}\widetilde{x_1^j} = \int_0^1 \partial_{x_t^j}\mathcal{E} d_{x_t^j}\phi_{t1}\widetilde{v}_t(x_t^j)dt \\ &= \int_0^1 \langle k_V(x_t^j,\cdot) d_{x_t^j}\phi_{t1}^* \nabla_{x_t^j}\mathcal{E}, \widetilde{v}_t\rangle_V dt\end{aligned}$$

We have reduced the computation to the derivative of $|\mathcal{C}(S_1) - \mathcal{C}(T)|^2_{W^*}$ with respect to the vertices of S_1. Let $\partial_{f_1^i}\mathcal{E}$ denote the contribution of a face f_1 to its vertex f_1^i. We have $\partial_{f_1^i}\mathcal{E}\ \eta = 2[\partial_{f_1^i}\mathcal{C}(S_1)\ \eta](\omega)$, where $\omega = K_W(\mathcal{C}(S_1) - \mathcal{C}(T))$ and $\mathcal{C}(S_1)(\omega) = \sum_f \delta^{N(f_1)}_{c(f_1)}(\omega) = \sum_f N(f_1)\cdot\overline{\omega}(c(f_1))$. Thus,

$$\begin{aligned}\partial_{f_1^i}\mathcal{E}\ \eta &= 2\partial_{f_1^i}N(f_1)\ \eta\cdot\overline{\omega}(c(f_1)) + N(f_1)\cdot d\overline{\omega}(c(f_1))\partial_{f_1^i}c(f_1)\ \eta, \\ &= (\eta\times e_1^i)\cdot\overline{\omega}(c(f_1)) + \frac{2}{3}N(f_1)\cdot d\overline{\omega}(c(f_1))\ \eta,\end{aligned}$$

so that $\nabla_{f_1^i}\mathcal{E} = (e_1^i\times\overline{\omega}(c(f_1))) + \frac{2}{3}d\overline{\omega}(c(f_1))^* N(f_1)$. It is left to compute $d_{x_t^i}\phi_{t1}$. It follows from properties of the differential that $\frac{d}{dt}(d_{x_t^i}\phi_{t1}) = -d_{x_t^i}\phi_{t1}d_{x_t^i}v_t$ (see [10]). Therefore we have the ordinary differential equation

$$\frac{d}{dt}\nabla_{x_t^i}\mathcal{E} = -(d_{x_t^i}v_t)^*\nabla_{x_t^i}\mathcal{E}, \tag{8}$$

which can be solved by integrating backward from time $t = 1$, since we can compute $d_{x_t^i}v_t$ using (7). Finally, we obtain $\nabla_{x_1^i}\mathcal{E}$ by summing $\nabla_{f_1^j}\mathcal{E}$ over all faces which share x_1^i as a vertex.

Gradient of J By a direct computation, the gradient of the regularization term, $\int_0^1 |v_t|_V^2$, is simply $2v_t$ so that the gradient of the functional J becomes

$$\nabla J_t(x) = 2\sum_j k_V(x_t^j, x)(d_{x_t^j}\phi_{t1}^*\nabla_{x_t^j}\mathcal{E} + \alpha_t^j) \tag{9}$$

3.4 Description of the algorithm

On the basis of remarks made in 3.3, we compute the functional and gradient as functions of the momentum variables α_t^j. The trajectories $x_t^j = \phi_t(x^j)$ being computed by solving the flow equation, written $\partial_t x_t^j = \sum_k k_V(x_t^k, x_t^j)\alpha_t^k$. Therefore the dimension of the parameter space is $3 * nt * nf$ where nt is the number of time steps and nf the number of faces of S. Equipped with equations (8) and (9), we implement a simple steepest descent algorithm.

Fig. 1. Left: template, right: target, center: mapped template.

4 Experiments

4.1 Experiments with faces dataset

In this experiment we used 10 segmented surfaces from the USF HumanID database [20] together with manually selected landmarks. The landmarks are used only for validation purposes. The first face was chosen as the template S to be matched to the other 9 surfaces. For each experiment we downsampled the original surfaces from 60 thousands triangles to 5 thousands triangles and we computed the optimal deformation between the downsampled meshes. Figure 1-left shows the image of the *original* template surface with its landmarks overlayed. Figure 1-right shows a target surface with

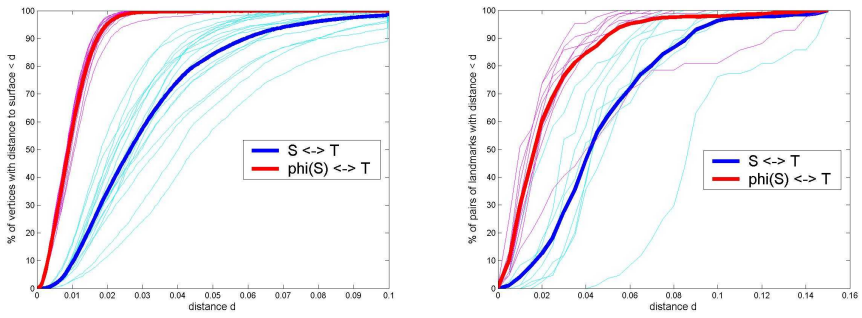

Fig. 2. Face experiments: distance graphs for vertices (left) and landmarks (right). Mean distance is given by the bold curves

its landmarks as well as the landmarks of the mapped surface. The mapped template is shown in the center panel. Figure 2-left shows distance error graphs, before and after the matching process (i.e. between S and T and between $\phi(S)$ and T). The left graph plots

the percentage of vertices whose distance to the other surface is less than d, as a function of distance d. On the right we plot the distance graphs for the sets of landmarks, i.e. the percentage of landmarks on target T such that the distance to their corresponding landmarks on S (resp. $\phi(S)$) is less than d. Note that the matching is visually satisfying, which is confirmed by the surface distance graph. However, success in matching some of the landmarks such as those along the chin, neck and jaw was not achieved. In fact, this may reflect the somewhat unreliable choice of these landmark points, which do not necessarily correspond to clearly defined features that can be reliably identified.

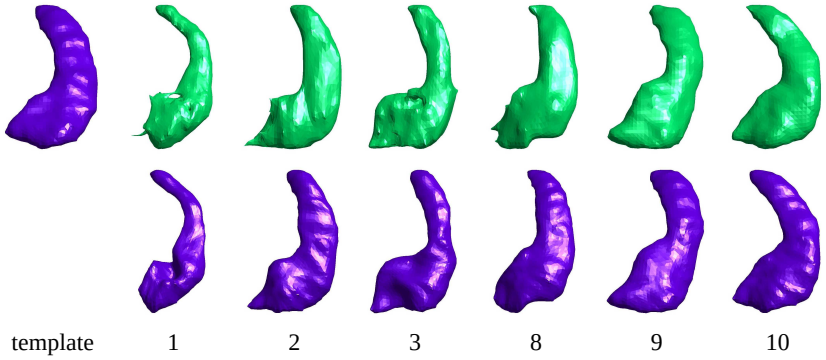

Fig. 3. up: 7 left hippocampi segmented surfaces (1,2 and 3 are of Alzheimer type). bottom: deformations of template through the action of the optimal diffeomorphisms.

4.2 Experiments on hippocampus data

Next we applied the matching algorithm to 15 left hippocampi segmented surfaces (see [21] for the method used); the first 7 belong to patients with Alzheimer disease and the others belong to normal subjects. In this experiment, the surfaces were downsampled to 500 triangles.

Figure 3 displays the deformations of the template surface from the matching process. Figure 4.3 shows the distance graph for this set of experiments. Note that for almost all vertices the distance is lower than 2mm. Also notice the small variance for these experiments.

4.3 Experiments on planum temporale segmented surfaces

As a final experiment, we applied the surface matching algorithm to segmentations of left and right planum temporales (PT) from 17 different subjects with 8 having auditory disorders. There is a high variability in sizes and shapes for this part of the brain, even between normal subjects, and also between left and right PTs of the same subject.

We chose to run two types of experiments on this set of data. In the first experiment we fixed one PT surface as the template and then registered it to the other 16 PTs

with our current matching algorithm. This was done for both left and right sets of data. Bilateral PT asymmetry studies are an active area of research, so in the second experiment we mapped each PT to its symmetric pair of the same subject. I.e. for each left PT we used its corresponding right PT as the template, and conversely for each right PT we used its corresponding left PT as the template. Figure 5 shows the distance graphs obtained for the left PT data of the first experiment, and for the left to right symmetry matchings.

The mean performance for this data was quite good, and similar to the performance for the hypocampus. Note, however the high variance of the distance measures in the graphs. This may be explained by the fact that the boundaries of the mapped template and target need not match if the geometries near the boundary are quite different from one another. For example, in Figure 6 we show a template on the left and two targets in the top row. In the bottom row are the mapped templates corresponding to the target of the same column. Overlayed on the template are manually defined landmarks which are also flowed under the mapping and overlayed on the mapped template. The landmarks shown on the target were estimated from the mapped template by choosing the closest point on the target for each landmark on the mapped template. In the first case (left) the algorithm gives correspondences which are consistent with what one may select by hand, whereas in the second case it gives correspondences for c and d landmarks which are not on the bound-

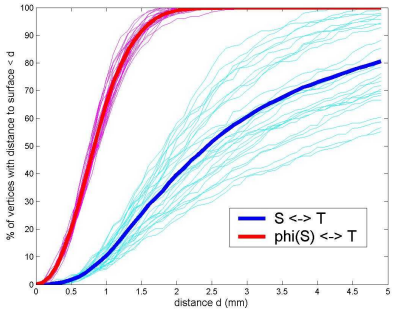

Fig. 4. Distance graph for the hippocampi experiments

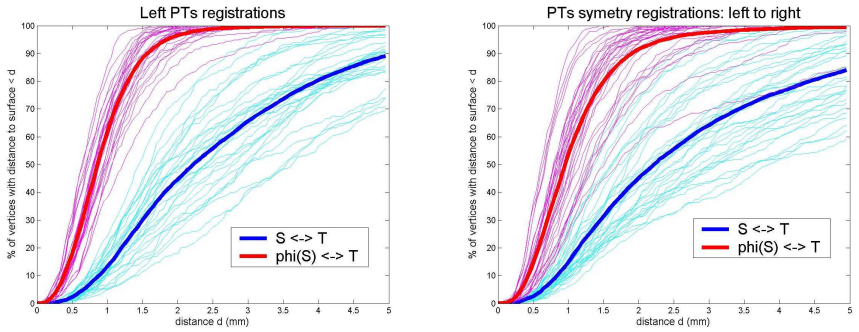

Fig. 5. Distance graphs for the planum temporale experiments.

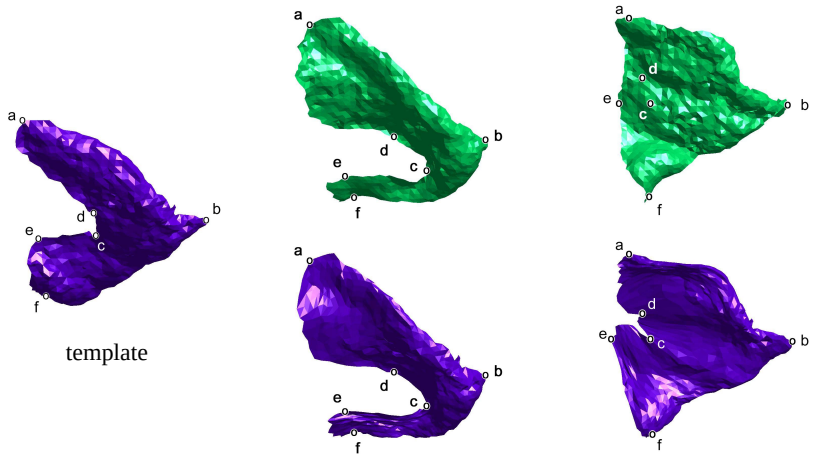

Fig. 6. Planum temporale experiments: correspondences of landmarks selected on template.

ary. Indeed here there are no obvious corresponding landmarks on the target since its shape is globally different from template.

5 Conclusion

We have presented a novel matching criterion for surfaces based on a sound theoretical framework. We integrated this criterion into a large deformation based variational problem, derived a discrete version and an algorithm for implementing its optimization via gradient descent. Finally we demonstrated its performance on different types of data. The main contribution was in recognizing currents as an appropriate mathematical modeling object for surfaces. Given the recent active developments in exterior calculus [22] and current based approaches to curvature estimation [23], we expect the representations to become more sophisticated (perhaps incorporating second order geometric information), and that discretization will continue to get better.

A promising and exciting immediate application of the diffeomorphic matching in this paper is in statistical inference of shape via momentum representation of flow, as described in [18]. It has been shown in [17] that the image of the template S under the flow ϕ_t is completely determined by the momentum (α^i) at time $t = 0$. Hence, the momenta encode the non-linear transformation from one structure into another, and furthermore, they live in a linear space which lends itself to linear statistical analysis.

6 Acknowledgements

Hippocampus surfaces were provided by M. Chupin, Cognitive Neuroscience and Brain Imaging Laboratory, CNRS, Paris and B. Dubois, Neurology Unit, Hopital de La Salpetriere, Paris. This research was partially supported by NIH grants 5R01MH064838-02, 5P41RR015241-04 and 2P01AG003991-21.

References

1. D'Arcy W. Thompson. *On Growth and Forms*. Cambridge University Press, Cambridge, England, 1917.
2. L Bookstein, F. *Morphometric tools for landmark data; geometry and biology*. Cambridge University press, 1991.
3. C. Davatzikos. Spatial transformation and registration of brain images using elastically deformable models. *Comp. Vision and Image Understanding*, 66(2):207–222, May 1997.
4. S. C. Joshi and M. I. Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Trans. Image Processing*, 9(8):1357–1370, 2000.
5. V. Camion and L. Younes. Geodesic interpolating splines. *EMMCVPR*, pages 513–527, 2001.
6. J. Glaunès, M. Vaillant, and M. I. Miller. Landmark matching via large deformation diffeomorphisms on the sphere. *Journal of Mathematical Imaging and Vision, MIA 2002 special issue*, 20, 2004.
7. H. Chui and A. Rangarajan. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89:114–141, 2003.
8. Yongmei Wang, Bradley S. Peterson, and Lawrence H. Staib. 3d brain surface matching based on geodesics and local geometry. *Computer Vision and Image Understanding*, 89:252–271, 2003.
9. R. H. Davies, T. F. Cootes, and C. J. Taylor. 3d statistical shape models using direct optimisation of description length. In *ECCV*, 2002.
10. J. Glaunès, A. Trouvé, and L. Younes. Diffeomorphic matching of distributions: A new approach for unlabelled point-sets and sub-manifolds matching. In *CVPR*, pages 712–718. IEEE Computer Society, 2004.
11. G. deRham. Variétés différentiables, formes, courants, formes harmoniques. *Act. Sci. Indust.*, 1222, 1955.
12. F. Morgan. *Geometric measure theory, 2nd ed.* Acad. Press, INC., 1995.
13. M.P. do Carmo. *Differential Forms and Applications*. Springer-Verlag, 1994.
14. G. Wahba. *Spline Models for Observational Data*. CBMS-NSF Regional conference series. SIAM, 1990.
15. A. Trouvé. An infinite dimensional group approach for physics based models. Technical report (electronically available at http://www.cis.jhu.edu), 1995.
16. P. Dupuis, U. Grenander, and M. I. Miller. Variational problems on flows of diffeomorphisms for image matching. *Quaterly of Applied Math.*, 56:587–600, 1998.
17. M. I. Miller, A. Trouvé, and L. Younes. Geodesic shooting in computational anatomy. Technical report, Center for Imaging Science, Johns Hopkins University, 2003.
18. M. Vaillant, M. I. Miller, L. Younes, and A. Trouvé. Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, 23:161–169, 2004.
19. M. I. Miller, A. Trouvé, and L. Younes. On the metrics and Euler-Lagrange equations of computational anatomy. *Annual Review of Biomedical Engineering*, 4:375–405, 2002.
20. USF HumanID 3D faces database, courtesy of Professor Sudeep Sarkar, University of South Florida, Tampa FL. Http://marthon.csee.usf.edu/HumanID/.
21. M. Chupin, D. Hasboun, S.Baillet, S. Kinkingnhun, B.Dubois, and L. Garnero. Competitive segmentation of the hippocampus and the volumetry in alzheimer's disease. In *10th Annual Meeting of the Organization for Human Brain Mapping*, June 13-17,2004.
22. A.N. Hirani. *Discrete exterior calculus.* PhD thesis, California Institute of Technology, 2003.
23. David Cohen-Steiner and Jean-Marie Morvan. Restricted delaunay triangulations and normal cycle. In *SCG '03: Proceedings of the nineteenth annual symposium on Computational geometry*, pages 312–321. ACM Press, 2003.