International Journal of Computer Vision 53(2), 153–167, 2003

Deformotion: Deforming Motion, Shape Average and the Joint Registration and Approximation of Structures in Images

ANTHONY J. YEZZI
Georgia Institute of Technology, Atlanta, GA 30332, USA
ayezzi@ece.gatech.edu

STEFANO SOATTO
University of California, Los Angeles, CA 90095, USA
soatto@ucla.edu

Received December 3, 2001; Revised January 30, 2003; Accepted January 30, 2003

Abstract. What does it mean for a deforming object to be "moving"? How can we separate the overall motion (a finite-dimensional group action) from the more general deformation (a diffeomorphism)? In this paper we propose a definition of motion for a deforming object and introduce a notion of "shape average" as the entity that separates the motion from the deformation. Our definition allows us to derive novel and efficient algorithms to register non-identical shapes using region-based methods, and to simultaneously approximate and align structures in greyscale images. We also extend the notion of shape average to that of a "moving average" in order to track moving and deforming objects through time. The algorithms we propose extend prior work on landmark-based matching to smooth curves, and involve the numerical integration of partial differential equations, which we address within the framework of level set methods.

Keywords: shape, segmentation, motion, registration

1. Introduction

Consider a sheet of paper falling. If it were a rigid object, one could describe its *motion* by providing the coordinates of one particle and the orientation of an orthogonal reference frame attached to that particle. That is, 6 numbers would be sufficient to describe the object at any instant of time. However, being a non-rigid object, in order to describe it at any instant of time one should really specify the trajectory of each individual particle on the sheet (Arnold, 1978). That is, if γ_0 represents the initial collection of particles, one could provide a function f that describes how the entire set of particles evolves in time: $\gamma_t = f(\gamma_0, t)$. Indeed, if each particle can move independently, there may be no notion of "overall motion," and a more appropriate description of f is that of a "*deformation*" of the sheet. That includes as a special case a rigid motion, described collectively by a rotation matrix $R(t) \in SO(3)$ and a translation vector $T(t) \in \mathbb{R}^3$, so that $\gamma_t = f(\gamma_0, t) = R(t)\gamma_0 + T(t)$ with $R(t)$ and $T(t)$ independent of the particle in γ_0. In practice, however, that is *not* how one usually describes a sheet of paper falling. Instead, one may say that the sheet is "moving" downwards along the vertical direction while "deforming." That is, even when the object is not rigid, one may still want to retain a notion of overall, or "global," motion, and describe departures from rigidity as a "deformation." This stems from one's desire to capture the fact that the sheet of paper is somehow moving as a whole, and its particles do not just behave like a swarm of bees. The jellyfish in Fig. 1 is just another example to illustrates the same

Figure 1. A jellyfish is "*moving while deforming*." What exactly does it mean? How can we separate its "global" motion from its "local" deformation?

issue. It certainly "moves," and it certainly "deforms" in the process of moving.

But what does it even mean for a deforming object to be "moving"? From a mathematical standpoint, rigorously defining a notion of motion for deforming objects presents a challenge. In fact, if we describe the deformation f as the composition of a rigid motion $(R(t), T(t))$ and a "deformation" function $h(\cdot, t)$, so that $\gamma_t = h(R(t)\gamma_0 + T(t), t)$, we can always find infinitely many different choices $\tilde{h}(\cdot, t)$, $\tilde{R}(t)$, $\tilde{T}(t)$ that give rise to the same overall deformation f:

$$\gamma_t = f(\gamma_0, t) = h(R(t)\gamma_0 + T(t), t) = \tilde{h}(\tilde{R}\gamma_0 + \tilde{T}(t), t)$$

by simply choosing $\tilde{h}(\gamma, t) \doteq h(R\tilde{R}^T(\gamma - \tilde{T}) + T, t)$ for any rigid motion $(\tilde{R}, \tilde{T})$. Therefore, we could describe the motion of our sheet with (R, T) as well as with $(\tilde{R}, \tilde{T})$, which is arbitrary, and in the end we would have failed in defining a notion of "motion" that is unique to the event observed.

So, how can we define a notion of motion for a deforming object in a mathematically sound way that reflects our intuition? The relevance of this problem goes beyond describing a falling sheet of paper. For instance, in Fig. 6, how do we describe the "motion" of a jellyfish? Or in Fig. 5 the "motion" of a storm? In neuroanatomy, how can we "register" a database of images of a given structure, say the corpus callosum (Fig. 9), by "moving" them to a common reference frame? In a defense scenario, how can we "track" targets that deform as they move, for instance a tank with a rotating turret?

All these questions ultimately boil down to an attempt to *separate the overall motion from the more general deformation.* Before proceeding, note that this is not always possible or even meaningful. In order to talk about the "motion" of an object, one must assume that "*something*" of the object is preserved as it deforms. For instance, it may not make sense to try to capture the "motion" of a swarm of bees, or of a collection of particles that indeed all move independently. What we want to capture mathematically is the notion of overall motion when indeed there is one that corresponds to our intuition!

The key to this paper is the observation that the notion of motion and the notion of shape are very tightly coupled. Indeed, we will see that our definition of *shape average is exactly what allows separating the motion from the deformation*. As a consequence, in our framework the process of "registering" a collection of shapes provides automatically an estimate of their *average*. Similarly, the process of segmenting a collection of images naturally results in their automatic alignment.

We now proceed to make the discussion above precise in a formal setting. We first propose our definitions for the simplest case where the "object" is a one-dimensional contour in Section 2, and later extend it to more general objects and more general notions of motion. We then give a detailed derivation of an algorithm to compute shape and motion in Section 3, which also results in an efficient way to compute the distance between planar shapes. When an object is being tracked over time, the notion of shape average is extended to that of a "moving average" (Section 4). We then extend these results from geometric shapes to images, resulting in their simultaneous approximation and registration in Section 5. Finally, in Section 6, we show results on a representative set of synthetic shapes as well as on real image sequences that illustrate our theory.

Before all that, in the next two sections we give a succinct description of the vast literature on shape and motion and how it relates to the contributions of our research.

1.1. Prior Related Work

The study of shape spans at least a hundred years of research in different communities from mathematical morphology to statistics, geology, neuroanatomy, paleontology, astronomy etc. Some of the earlier attempts to formalize a notion of shape include D'Arcy Thompson's treatise "Growth and Form" (Thompson,

1917), the work of Matheron on "Stochastic Sets" (Matheron, 1975) as well as that of Thom (1975), Giblin (1977) and others.

In statistics, the study of "Shape Spaces" was championed by Kendall (1984), Le and Kendall (1993), Mardia and Dryden (1989), and Carne (1990). Shapes are defined as the equivalence classes of N *points* in $\mathbb{R}^M$ under the similarity group, $\mathbb{R}^{MN}/\{SE(M) \times \mathbb{R}\}$. Shape spaces are thus organized in a fiber bundle where motion along the fibers corresponds to rotated, translated and scaled versions of the collection of points, while motion across fibers corresponds to deformations that change their mutual position individually (not just by a global scale). The fiber bundle can be endowed with a metric structure and with probability measures that allow comparing the so-defined shapes and compute statistics of collections of shapes. These tools have proven useful in contexts where distinct "*landmarks*" are available, for instance in comparing biological shapes with N distinct "parts." However, comparing objects that have a different number of parts, or objects that do not have any distinct landmark, is elusive within the framework of statistical shape spaces. Although the framework clearly distinguishes the notion of "motion" (along the fibers) from the "deformation" (across fibers), the analytical tools are essentially tied to the point-wise representation. One of our goals in this paper is to extend the theory to smooth curves,[1] surfaces and other geometric objects that do not have distinct "landmarks."

In computer vision, a wide literature exists for the problem of "matching" or "aligning" objects based on their images, and space limitations do not allow us to do justice to the many valuable contributions. We refer the reader to Veltkamp and Hagedoorn (1999) for a recent survey. A common approach consists of matching collections points organized in graphs or trees (e.g. Lades et al., 1993; Fischler and Elschlager, 1973). Belongie, et al. (2001) propose comparing planar contours based on their "shape context." There, points are not bound to represent particular "landmarks" but are just a discrete representation of the contour. Their matching is, by construction, invariant with respect to either the affine or the Euclidean group, and the resulting match is based on "features" rather than on image intensity directly, similarly to Chui and Rangarajan (2000) and Dutta and Jain (2001). Koenderink (1990) is credited with providing some of the key ideas involved in formalizing a notion of shape that matches our intuition. However, Mumford has critiqued current theories of shape on the grounds that they fail to capture the essential features of perception (Mumford, 1991).

"Deformable Templates," pioneered by Grenander (1993), do not rely on "features" or "landmarks;" rather, images are directly deformed by a (possibly infinite-dimensional) group action and compared for the best match in an "image-based" approach (Yuille, 1991). There, the notion of "motion" (or "alignment" or "registration") coincides with that of deformation, and there is no clear distinction between the two (Bereziat et al., 1997). Grenander's work sparked a current that has been particularly successful in the analysis of medical images, for instance (Grenander and Miller, 1994). We would like to retain some of the power and flexibility of deformable templates, but within this framework mark a clear distinction between "motion" and "deformation."

Another line of work uses variational methods and the solution of partial differential equations (PDEs) to model shape and to compute distances and similarity. In this framework, not only can the notion of alignment or distance be made precise (Azencott et al., 1996; Younes, 1998; Miller and Younes, 1999; Kimmel and Bruckstein, 1995; Samson et al., 1999), but quite sophisticated theories that encompass perceptually relevant aspects, can be formalized in terms of the properties of the evolution of PDEs (e.g. Kimmel et al., 1998). The work of Kimia et al. (1995) and Sebastian et al. (2000) describes a scale-space that corresponds to various stages of evolution of a diffusing PDE, and a "reacting" PDE that splits "salient parts" of planar contours by generating singularities. Kimia et al. (1995) also contains a nice taxonomy of existing work on shape and deformation and a review of the state of the art as of 1994.

The variational framework has also proven very effective in the analysis of medical images (Malladi et al., 1995, 1996; Thompson and Toga, 1996). Although most of the ideas are develop in a deterministic setting, many can be transposed to a probabilistic context (e.g. Zhu et al., 1995). None of these approaches, however, distinguishes a notion of motion that is separate from the deformation; the evolution of shapes is driven by energy and regularization terms, rather than by the action of a finite-dimensional group of transformations. We would like to extend this framework to evolve contours simultaneously with respect to a group element and a generic deformation, and try to infer both from data and render them separate or "independent" in a precise sense.

A common approach to matching planar contours within the context of scale-space is to *not* match the contours directly, but to first represent them through a common scale-space and then match a given scale, or even all scales. The rationale being that, even if the original contours are not well matched by a group action, their representations through scale-space at some scale may be. Scale-space is a very active research area, and some of the key contributions as they relate to the material of this paper can be found in Jackway and Deriche (1996), ter Haar Romeny et al. (1997), Kimmel (1997), Alvarez and Morel (1994), and Alvarez et al. (1993, 1999) and references therein. The "alignment," or "registration," of curves has also been used to define a notion of "shape average" by several authors (see Leventon et al., 2000 and references therein). The shape average, or "prototype," can then be used for recognition in a nearest-neighbor classification framework, or to initialize image-based segmentation by providing a "prior." Leventon et al. (2000) perform principal component analysis in the aligned frames to regularize the segmentation of regions with low contrast in brain images. However, the alignment is performed ad-hoc by pre-processing the images, rather than posing it as part of the inference problem. Errors in the pre-processing stage can never be compensated. Similarly, Yezzi and Soatto (2001) performs the joint segmentation of a number of images by assuming that their registration (stereo calibration) is given. We wish to extend these approaches to situations where the calibration/registration is not known a-priori. A somewhat complementary work is Yezzi et al. (2001), where objects, assumed to be identical except for a group action, are registered by minimizing a region-based cost functional. We wish to extend that approach to situations where the objects are not "similar" (i.e. equivalent under the group action), but they also undergo deformations.

Also related to this paper is the recent work of Paragios and Deriche, where active regions are tracked as they "move." In Paragios and Deriche (2000) the notion of motion is not made distinct from the general deformation, and therefore what is being tracked is a general (infinite-dimensional) deformation. Our aim is to define tracking as a trajectory on a finite-dimensional group, despite infinite-dimensional deformations. Substantially different in methods, but related in the intent, is the work on stochastic filters for contour tracking and snakes (see Blake and Isard, 1998 and references therein). There, however, what is being tracked over time is a general deformation (although finitely parametrized via splines or other parametric descriptions), rather than a (group) motion. Therefore, the end product of these tracking algorithms is not a trajectory on a finite-dimensional group, but a generic sequence of deformations.

1.2. Contributions of This Paper

We wish to warn the reader at the outset that we do *not* intend to present a comprehensive theory of shape that captures the complexity and intricacy of the problem or that subsumes and generalizes existing theories. Rather, within this vast theme, we have identified the particular issue of "separating" the notion of motion from a more general deformation as a crucial one, on which we wish to say something fairly precise. The consequences of our analysis are robust algorithms for matching, registering and tracking deforming objects, computing a meaningful notion of "shape average" and the distance between shapes.

The situations we wish to describe are those where *objects undergo a distinct overall "global" motion while "locally" deforming*.[2] Our approach does not apply when objects deform wildly, when different "parts" of the object undergo different deformations, and it entails no notion of hierarchy or compositionality.

Under these assumptions, our contribution consists of (1) a novel definition of motion for a deforming object and (2) a corresponding definition of shape average (Section 2). Our definition allows us to derive novel and efficient algorithms to (3) register non-identical (or non-equivalent) shapes using region-based methods. We use our algorithms to (4) simultaneously approximate and register structures in images, or to simultaneously segment and calibrate images (Section 5). In the context of tracking, we extend our definition to a novel notion of (5) "moving average" of shape, and use it to (6) perform tracking for deforming objects (Section 4).

Our definition of motion and shape average does not rely on a particular representation of objects (e.g. explicit vs. implicit, parametric vs. non-parametric), nor on the particular choice of group (e.g. affine, Euclidean), nor is it restricted to a particular modeling framework (e.g. deterministic, energy-based vs. probabilistic). For the implementation of our algorithms on deforming contours, we have chosen an implicit non-parametric representation in terms of level sets, following Osher and Sethian (1988), and we

have implemented numerical algorithms for integrating PDEs to converge to the steady-state of an energy-based functional. However, these choices can be easily changed without altering the nature of the contribution of this paper.

Naturally, since shape and motion are computed as the solution of a nonlinear optimization problem, the algorithms we propose are only guaranteed to converge to local minima and, in general, no conclusions can be drawn on uniqueness. Indeed, it is quite simple to generate pathological examples where the setup we have proposed fails. In the experimental section we will highlight the limitations of the approach when used beyond the assumptions for which it is designed.

2. Defining Motion and Shape Average

The key idea underlying our framework is that the notion of *motion* throughout a deformation is very tightly coupled with the notion of *shape average*. In particular, if a deforming object is recognized as moving, there must be an underlying object (which will turn out to be the shape average) moving with the same motion, from which the original object can be obtained with minimum deformations. Therefore, we will model a general deformation as the composition of a group action g on a particular object, on top of which a local deformation is applied. The shape average is defined as the one that minimizes such deformations.

Let $\gamma_1, \gamma_2, \ldots, \gamma_n$ be n "shapes" (we will soon make the notion precise.) Let the mappings between each pair of shapes be T_{ij}

$$\gamma_i = T_{ij}\gamma_j, \quad i, j = 1 \ldots n. \tag{1}$$

Each comprises the action of a group $g \in G$ (e.g. $G = SE(2)$) and a pair more general transformations h_i, h_j that belong to a pre-defined class $\mathcal{H}$ (for instance diffeomorphisms). The deformations are not arbitrary, but originate from a common "shape" μ, defined in such a way that

$$\gamma_i = h_i \circ g_i(\mu), \quad i = 1 \ldots n. \tag{2}$$

Therefore, in general, following the commutative diagram of Fig. 2, we have that

$$T_{ij} \doteq h_i \circ g_i \circ g_j^{-1}(\mu) \circ h_j^{-1} \tag{3}$$

so that $g = g_i g_j^{-1}$. Given two or more "shapes" and a cost functional $E : \mathcal{H} \to \mathbb{R}^+$ defined on the set of diffeomorphisms, the motion g_t and the shape average are defined as the minimizers of $\sum_{i=1}^{n} E(h_i)$ subject to $\gamma_i = h_i \circ g_i(\mu)$. Since all that matter in the cost of h_i are the "shapes" before and after the transformation, $\mu_i \doteq g_i(\mu)$ and γ_i, we can write, with an abuse of notation, $E(h(\mu_i, \gamma_i)) \doteq E(\mu_i, \gamma_i)$. We are therefore ready to define our notion of motion during a deformation.

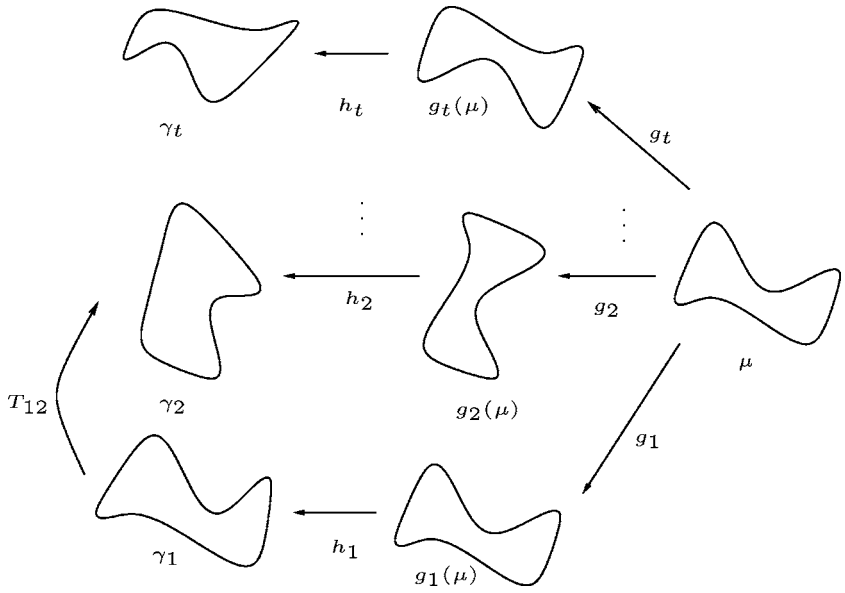

Figure 2. A model (commutative diagram) of a deforming contour.

Definition 1. Let $\gamma_1, \ldots, \gamma_n$ be a collection of compact hypersurfaces embedded in $\mathbb{R}^N$ (in this paper we concentrate on $N = 3$), which we call pre-shapes. Let $\mathcal{H}$ be a class of diffeomorphisms acting on γ_i, and let $E : \mathcal{H} \to \mathbb{R}^+$ be a positive, real-valued functional. Consider now a group G acting on γ_i via $g(\gamma_i)$. We say that $\hat{g}_1, \ldots, \hat{g}_n$ is a **motion** undergone by γ_i, $i = 1 \ldots n$, if there exists a pre-shape $\hat{\mu}$ such that

$$\boxed{\begin{array}{l} \hat{g}_1, \ldots, \hat{g}_n, \hat{\mu} = \arg\min_{g_i, \mu} \sum_{i=1}^{n} E(h_i) \\ \text{subject to } \gamma_i = h_i \circ g_i(\mu) \quad i = 1 \ldots n \end{array}} \tag{4}$$

The pre-shape $\hat{\mu}$ is called the **shape average** relative to the group G, or G-average, and the quantity $\hat{g}_i^{-1}(\gamma_i)$ is called the **shape** of γ_i.

Remark 1 (Invariance). In the definition above, one will notice that the shape average is actually a pre-shape, and that there is an arbitrary choice of group action g_0 that, if applied to γ_i and μ, leaves the definition unchanged (the functional E is invariant with respect to g_0 because $T_{ij}(g \circ g_0, h_i \circ g_0, h_j \circ g_0) = T_{ij}(g, h_i, h_j) \, \forall \, g_0$). For the case of the Euclidean group $SE(N)$, a way to see this is to notice that the reference frame where μ is described is arbitrary. Therefore, one may choose, for instance, $\mu = h_1^{-1}(\gamma_1)$.

Remark 2 (G-average). Notice that the notion of *shape average* above is relative to the particular choice of group G. For instance, given a number of pre-shapes $\gamma_1, \ldots, \gamma_n$, the shape average $\hat{\mu}$ relative to the Euclidean group will, in general, be different than the shape average relative to the affine or the projective group (e.g. Fig. 3). Therefore, when we talk about average, we always have to specify the group G, e.g. Euclidean average, affine average etc.

Remark 3 (Symmetries). In Definition 1 we have purposefully avoided to use the article "*the*" for the minimizing value of the group action $\hat{g}_i$. It is in fact possible that the minimum of (4) may not be unique. A particular case when this occurs is *when the pre-shape* γ is (symmetric, or) invariant with respect to an entire subgroup of G. Another way to say this is that the pre-shape is *symmetric* with respect to a subgroup of G. For instance, consider the set of closed contours described in the previous section, and let g_0 be such that $g_0(\gamma) = \gamma \; \forall \; g_0 \in G_0 \subset G$. Then, clearly $\hat{g}$ and $\hat{g} \circ g_0$ produce the same value in the functional E, and therefore the two are *indistinguishable* from the data. The simplest case is a circular contour, which is invariant with respect to rotations around its center. It is clear that by matching circles we can detemine the relative position of their centers, but not the relative orientation of the reference frame attached to each circle, since the latter is arbitrary. Notice, however, that the notion of *shape average* is still well-defined even when the notion of motion is not unique. This is because any element in the symmetry group suffices to register the pre-shapes, and therefore compute the shape average (Fig. 3).

In Section 3 we specialize this definition for the case of a planar contour undergoing Euclidean or affine motion and differentiable deformations, and we show how to compute motion, shape average, as well as distances between shapes.

3. Shape and Deformation of a Planar Contour

In this section we consider the implementation of the program above for a simple case: two closed planar contours, γ_1 and γ_2, where we choose as the cost functional for the deformations h_1, h_2 either the set-symmetric difference Δ of their interior (the union minus the intersection of μ_i and $h_i(\mu_i)$), or what we call the signed distance score[3] ψ

$$\psi(\mu_i, \gamma_i) \doteq \int_{\bar{\mu}_i} \zeta(\gamma_i) \, d\mathbf{x} \tag{5}$$

where $\bar{\mu}_i$ denotes the interior of the contour μ_i and ζ is the signed distance function of the contour γ_i; $d\mathbf{x}$ is the area form on the plane. In either case, since we have an arbitrary choice of the global reference frame, we can choose $g_1 = e$, the group identity. We also call $g \doteq g_2$, so that $\mu_2 = g(\mu)$. The problem of defining the motion and shape average can then be written as

$$\begin{array}{c} \hat{g}, \hat{\mu} = \arg\min_{g, \mu} \sum_{i=1}^{2} E(h_i) \\ \text{subject to } \gamma_1 = h_1(\mu); \; \gamma_2 = h_2 \circ g(\mu). \end{array} \tag{6}$$

As we have anticipated, we choose either $E(h_i) = \Delta(g_i(\mu), \gamma_i)$ or $E(h_i) \doteq \psi(g_i(\mu), \gamma_i)$. Therefore, abusing the notation as anticipated before Definition 1, we can write the problem above as an unconstrained

minimization

$$\boxed{\hat{g}, \hat{\mu} = \arg\min_{g,\mu} \phi(\gamma_1, \gamma_2)} \quad \text{where}$$
$$\boxed{\phi(\gamma_1, \gamma_2) \doteq E(\mu, \gamma_1) + E(g(\mu), \gamma_2)} \tag{7}$$

and E is either Δ or ψ. The estimate $\hat{g}$ defines the motion between γ_1 and γ_2, and the estimate $\hat{\mu}$ defines the average of the two contours.

If one thinks of contours and their interior, represented by a *characteristic function* χ, as a binary image, then the cost functionals above are just particular cases of a more general cost functional where each term is obtained by integrating a function inside and a function outside the contours

$$\boxed{\phi = \sum_{i=1}^{2} \int_{\bar{\mu}_{in}} f_{in}(\mathbf{x}, \gamma_i)\, d\mathbf{x} + \int_{\bar{\mu}_{out}} f_{out}(\mathbf{x}, \gamma_i)\, d\mathbf{x}} \tag{8}$$

where the bar in $\bar{\mu}$ indicates that the integral is computed on a *region* inside or outside μ and we have emphasized the fact that the function f depends upon the contour γ_i. For instance, for the case of the set-symmetric difference, we have $f_{in} = (\chi_\gamma - 1)$ and $f_{out} = \chi_\gamma$. To solve the problem, therefore, we need to minimize the following functional

$$\int_{\bar{\mu}_{in}} f_{in}(\mathbf{x}, \gamma_1)\, d\mathbf{x} + \int_{\bar{\mu}_{out}} f_{out}(\mathbf{x}, \gamma_1)\, d\mathbf{x} + \int_{g(\bar{\mu}_{in})} f_{in}(\mathbf{x}, \gamma_2)\, d\mathbf{x} + \int_{g(\bar{\mu}_{out})} f_{out}(\mathbf{x}, \gamma_2)\, d\mathbf{x} \tag{9}$$

which can be written, after a change of variable in the second two terms and some rearranging, as

$$\int_{\bar{\mu}_{in}} f_{in}(\mathbf{x}, \gamma_1)\, d\mathbf{x} + \int_{\bar{\mu}_{out}} f_{out}(\mathbf{x}, \gamma_1)\, d\mathbf{x} + \int_{\bar{\mu}_{in}} f_{in}(g(\mathbf{x}), \gamma_2)|J_g|\, d\mathbf{x} + \int_{\bar{\mu}_{out}} f_{out}(g(\mathbf{x}), \gamma_2)|J_g|\, d\mathbf{x} \tag{10}$$

$$\int_{\bar{\mu}_{in}} f_{in}(\mathbf{x}, \gamma_1) + f_{in}(g(\mathbf{x}), \gamma_2)|J_g|\, d\mathbf{x} + \int_{\bar{\mu}_{out}} f_{out}(\mathbf{x}, \gamma_1) + f_{out}(g(\mathbf{x}), \gamma_2)|J_g|\, d\mathbf{x} \tag{11}$$

where $|J_g|$ is the determinant of the Jacobian of the group action g. This makes it easy to compute the component of the first variation of ϕ along the normal direction to the contour μ, so that we can impose

$$\nabla_\mu \phi \cdot N = 0 \tag{12}$$

to derive the first-order necessary condition. If we choose $G = SE(2)$, an isometry, it can be easily shown that (Zhu et al., 1995)

$$\boxed{\begin{aligned}\nabla_\mu \phi = {} & f_{in}(\mathbf{x}, \gamma_1) - f_{out}(\mathbf{x}, \gamma_1) \\ & + f_{in}(g(\mathbf{x}), \gamma_2) - f_{out}(g(\mathbf{x}), \gamma_2)\end{aligned}} \tag{13}$$

3.1. Representation of Motions

For the specific case of matrix Lie groups (e.g. $G = SE(2)$), there exist twist coordinates ξ that can be represented as a skew-symmetric matrix $\widehat{\xi}$ so that[4]

$$\boxed{g = e^{\widehat{\xi}}} \quad \text{and} \quad \boxed{\frac{\partial g}{\partial \xi_i} = \frac{\partial \widehat{\xi}}{\partial \xi_i} g} \tag{14}$$

where the matrix $\frac{\partial \widehat{\xi}}{\partial \xi_i}$ is composed of zeros and ones and the matrix exponential can be computed in closed form. In Appendix A we give the expression of the exponential for the case of $SO(2)$, $SE(2)$, $SO(3)$, $SE(3)$, known as Rodrigues' formula.

3.2. Variation with Respect to the Group Action

To compute the variation of the functional ϕ with respect to the group action g, we first notice that the first two terms in ϕ do not contribute since they are independent of g. Therefore, we are left with having to compute the variation of

$$\int_{g(\bar{\mu}_{in})} f_{in}(\mathbf{x}, \gamma_2)\, d\mathbf{x} + \int_{g(\bar{\mu}_{out})} f_{out}(\mathbf{x}, \gamma_2)\, d\mathbf{x}. \tag{15}$$

To simplify the derivation, we consider the case of $SE(3)$. Other cases follow along similar lines (except for the Jacobian of the transformations, which is one in the isometric case); we also note that both terms above are of the generic form $A(g) \doteq \int_{g(\bar{\mu})} f(\mathbf{x})\, d\mathbf{x}$. Therefore, we consider the variation of A with respect to the components of the twist ξ_i, $\frac{\partial A}{\partial \xi_i}$, which we will eventually use to compute the gradient with respect to the natural connection $\nabla_G \phi = \widehat{(\frac{\partial \phi}{\partial \xi})} g$. We first rewrite $A(g)$ using the change of measure

$\int_{g(\bar{\mu})} f(\mathbf{x})\,d\mathbf{x} = \int_{\bar{\mu}} f \circ g(\mathbf{x})|J_g|\,d\mathbf{x}$ which leads to $\frac{\partial A(g)}{\partial \xi_i} = \int_{\bar{\mu}} \frac{\partial}{\partial \xi_i}(f \circ g(\mathbf{x}))|J_g|d\mathbf{x} + \int_{\bar{\mu}} (f \circ g(\mathbf{x}))\frac{\partial}{\partial \xi_i}|J_g|d\mathbf{x}$ and note that the Euclidean group is an isometry and therefore the determinant of the Jacobian is one and the second integral is zero. The last equation can be re-written, using Green's theorem, as $\int_{g(\mu)} \langle f(\mathbf{x})\frac{\partial g}{\partial \xi_i} \circ g^{-1}(\mathbf{x}), N\rangle ds = \int_{\mu} \langle f \circ g(\mathbf{x})\frac{\partial g}{\partial \xi_i}, g_* N\rangle ds$ where g_* indicates the push-forward. Notice that g is an isometry and therefore it does not affect the arc length; we then have

$$\frac{\partial A(g)}{\partial \xi_i} = \int_{\mu} f(g(\mathbf{x}))\left\langle \frac{\partial \widehat{\xi}}{\partial \xi_i} g, g_* N \right\rangle ds \tag{16}$$

After collecting all the partial derivatives into an operator $\frac{\partial \phi}{\partial \xi}$, we can write the evolution of the group action.

3.3. Evolution

The algorithm for evolving the contour and the group action consists of a two-step process where an initial estimate of the contour $\hat{\mu} = \gamma_1$ is provided, along with an initial estimate of the motion $\hat{g} = e$, the identity of the group.[5] The contour and motion are then updated in an iterated minimization where motion is updated according to

$$\boxed{\frac{d\hat{g}}{dt} = \left(\widehat{\frac{\partial \phi}{\partial \xi}}\right)\hat{g}} \tag{17}$$

Notice that this is valid not just for $SE(2)$, but for any (finite-dimensional) matrix Lie group, although there may not be a closed-form solution for the exponential map like in the case of $SE(3)$ and its subgroups. In practice, the group evolution (17) can be implemented in local (exponential) coordinates by evolving ξ defined by $g = e^{\widehat{\xi}}$ via $\frac{d\xi}{dt} = \frac{\partial \phi}{\partial \xi}$. In the level set framework, the derivative of the cost function ϕ with respect to the coordinates of the group action ξ_i can be computed as the collection of two terms, one for f_{in}, one for f_{out} where $\frac{\partial \phi}{\partial \xi_i} = \int_{g(\gamma_{1,2})} \langle \frac{\partial g(\mathbf{x})}{\partial \xi_i}, f_{\{in,out\}}(g(\mathbf{x}), \gamma_{1,2}) J(g_* T)\rangle ds$. The contour $\hat{\mu}$ evolves according to

$$\boxed{\begin{aligned}\frac{d\hat{\mu}}{dt} &= (f_{in}(\mathbf{x}, \gamma_1) - f_{out}(\mathbf{x}, \gamma_1) \\ &\quad + f_{in}(g(\mathbf{x}), \gamma_2) - f_{out}(g(\mathbf{x}), \gamma_2))N.\end{aligned}} \tag{18}$$

As we have already pointed out, the derivation can be readily extended to surfaces in space.

3.4. Distance Between Shapes

The definition of motion $\hat{g}$ and shape average $\hat{\mu}$ as a minimizer of (7) suggests defining the distance[6] between two shapes as the "energy" necessary to deform one into the other via the average shape:

$$d(\gamma_i, \gamma_j) \doteq E(\gamma_i, T(\hat{g}, \hat{h})\gamma_j). \tag{19}$$

For instance, for the case of the set-symmetric difference of two contours, we have

$$\begin{aligned} d_{\Delta}(\gamma_1, \gamma_2) \doteq \int & \chi_{\hat{\mu}} + \chi_{\gamma_1} - 2\chi_{\hat{\mu}}\chi_{\gamma_1} + \chi_{\hat{g}(\hat{\mu})} \\ & + \chi_{\gamma_2} - 2\chi_{\hat{g}(\hat{\mu})}\chi_{\gamma_2}\,d\mathbf{x} \end{aligned} \tag{20}$$

and for the signed distance score we have

$$d_{\psi}(\gamma_1, \gamma_2) \doteq \int_{\hat{\tilde{\mu}}} \zeta(\gamma_1)\,d\mathbf{x} + \int_{\hat{g}(\hat{\tilde{\mu}})} \zeta(\gamma_2)\,d\mathbf{x}. \tag{21}$$

In either case, given two contours, a gradient flow algorithm based on Eqs. (17) and (18), when it converges to the global minimum, returns as the minimum value the distance between the shapes corresponding to the two contours.

4. Moving Average and Tracking

The discussion above assumes that an unsorted collection of shapes is available, where the deformation between any two shapes is "small" (modulo G), so that the whole collection can be described by a single average shape. Consider however the situation where an object is evolving in time, for instance Fig. 5. While the deformation between adjacent time instants could be captured by a group action and a small deformation, as time goes by the object may change so drastically that talking about a global time average may not make sense.

One way to approach this issue is by defining a notion of "*moving average*," similar to what is done in time series analysis.[7] In order to adapt this model to our case, the most significant change is the representation of the uncertainty during the evolution. In classical linear time series, uncertainty is modeled via additive noise. In our case, the uncertainty is an infinite-dimensional deformation h that acts on the measured contour. So the

model becomes

$$\begin{cases} \mu(t+1) = g(t)\mu(t) \\ \gamma(t) = h(\mu(t)) \end{cases} \tag{23}$$

where $\mu(t)$ represents the moving average of order $k = 1$. The procedure described in Section 3, initialized with $\mu(0) = \gamma_1$, provides an estimate of the moving average of order 1, as well as the *tracking* of the trajectory $g(t)$ in the group G, which in (23) is represented as the model parameter. Note that the procedure in Section 3 simultaneously estimates the state $\mu(t)$ and identifies the parameters $g(t)$ of the model (23). It does so, however, without imposing restrictions on the evolution of $g(t)$. If one wants to impose additional constraints on the motion parameters, one can augment the state of the model to include the parameters g.

$$\begin{cases} g(t+1) = e^{\widehat{\xi}(t)} g(t) \\ \mu(t+1) = g(t)\mu(t) \\ \gamma(t) = h(\mu(t)) \end{cases} \tag{24}$$

and specify restrictions on ξ. This, however, is beyond the scope of this paper. In Fig. 5 we show the results of tracking a storm with a moving average of order one.

5. Simultaneous Approximation and Registration of Non-Equivalent Shapes

So far we have assumed that the given shapes are obtained by moving and deforming a common underlying "template" (the average shape). Even though the given shapes are not *equivalent* (i.e. there is no group action g that maps one exactly onto the other), g is found as the one that minimizes the cost of the deviation from such an equivalence. In the algorithm proposed in Eqs. (17) and (18), however, there is no explicit requirement that the deformation between the given shapes be small. Therefore, the procedure outlined can be seen as an algorithm to register shapes that are not equivalent under the group action. A *registration* is a group element $\hat{g}$ that minimizes the cost functional (4).

To illustrate this fact, consider the two considerably different shapes shown in Fig. 7, γ_1, γ_2. The simultaneous estimation of their average μ, for instance relative to the affine group, and of the affine motions that best matches the shape average onto the original ones, g_1, g_2, provides a registration that maps γ_1 onto γ_2 and viceversa: $g = g_2 g_1^{-1}$.

In Fig. 9 we extend this approach to the simultaneous approximation and registration of a collection of 4 images of the corpus callosum. This is done by running a joint segmentation algorithm that simultaneously segments all images. Modulo smoothness constraints, this procedure converges in steady-state to the minimum squared-error thresholding of the collection of images. This is in general better than first segmenting images individually, and then running the procedure we have described in previous section to binary shapes.

6. Experiments

In this section we report the result of various experiments on both binary shapes and grayscale images. Since our procedure is based on a gradient descent, in principle it is subject to convergence to local minima. However, in all the experiments we have conducted, convergence is reached from a generic initial state (e.g. a circle or square on the plane, and a sphere in space).

Figure 3 illustrates the difference between the motion and shape average computed under the Euclidean group, and the affine one. The three examples show the two given shapes γ_i, the mean shape registered to the original shapes, $g_i(\mu)$ and the mean shape μ. Notice that affine registration allows us to simultaneously capture the square and the rectangle, whereas the Euclidean average cannot be registered to either one, and is therefore only an approximation.

Figure 4 compares the effect of choosing the signed distance score (left) and the set-symmetric difference (right) in the computation of the motion and average shape. The first choice results in an average that captures the common features of the original shapes, whereas the second captures more of the features in each one. Depending on the application, one may prefer one or the other.

Figure 5 shows the results of tracking a storm. The affine moving average is computed, and the resulting affine motion is displayed. The same is done for the jellyfish in Fig. 6.

Figures 7 and 8 are meant to challenge the assumptions underlying our method. The pairs of shapes chosen, in fact, are not simply local deformations of one another. Therefore, the notion of shape average is not meaningful *per se* in this context, but serves to compute

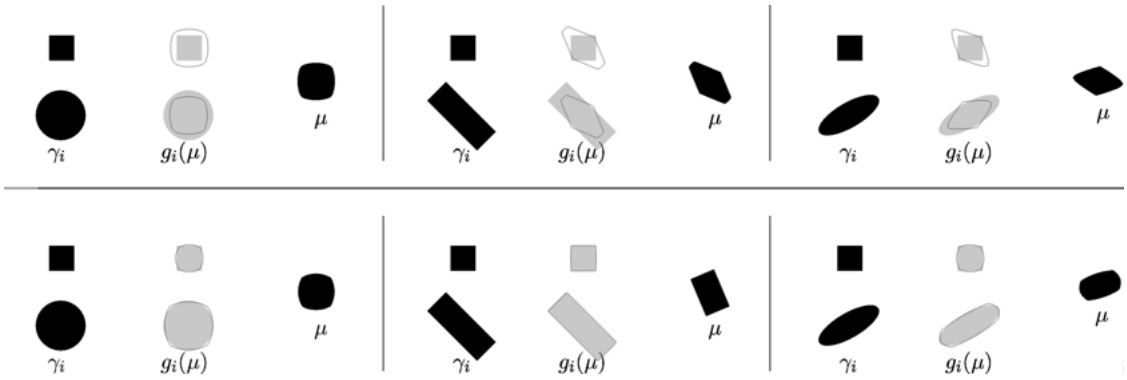

Figure 3. Euclidean (top) vs. affine (bottom) registration and average. For each pair of objects γ_1, γ_2, the registration $g_1(\mu)$, $g_2(\mu)$ relative to the Euclidean motion and affine motion is shown, together with the Euclidean average and affine average μ. Note that the affine average can simultaneously "explain" a square and a rectangle, whereas the Euclidean average cannot.

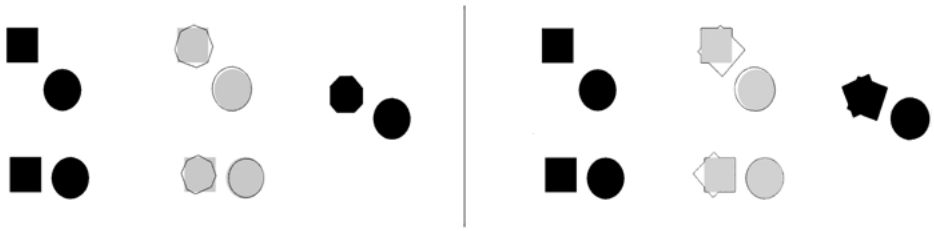

Figure 4. Signed distance score (left) vs. set-symmetric difference (right). Original contours (γ_1 on the top, γ_2 on the bottom), registered shape $g_i(\mu)$ and shape average μ. Note that the original objects are not connected, but are composed by a circle and a square. The choice of pseudo-distance between contours influences the resulting average. The signed distance score captures more of the features that are common to the two shapes, whereas the symmetric difference captures the features of both.

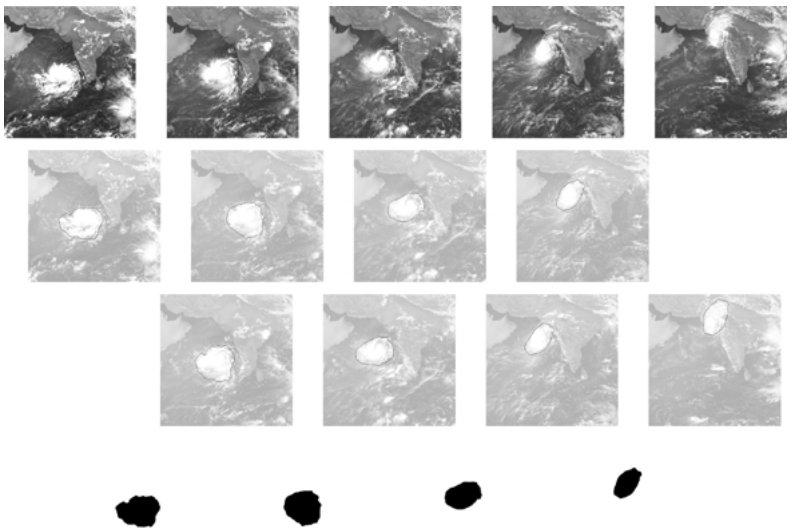

Figure 5. Storm (first row) a collection of images from EUMETSAT ©2001, affine motion of the storm based on two adjacent time instances, (bottom) moving average of order 1.

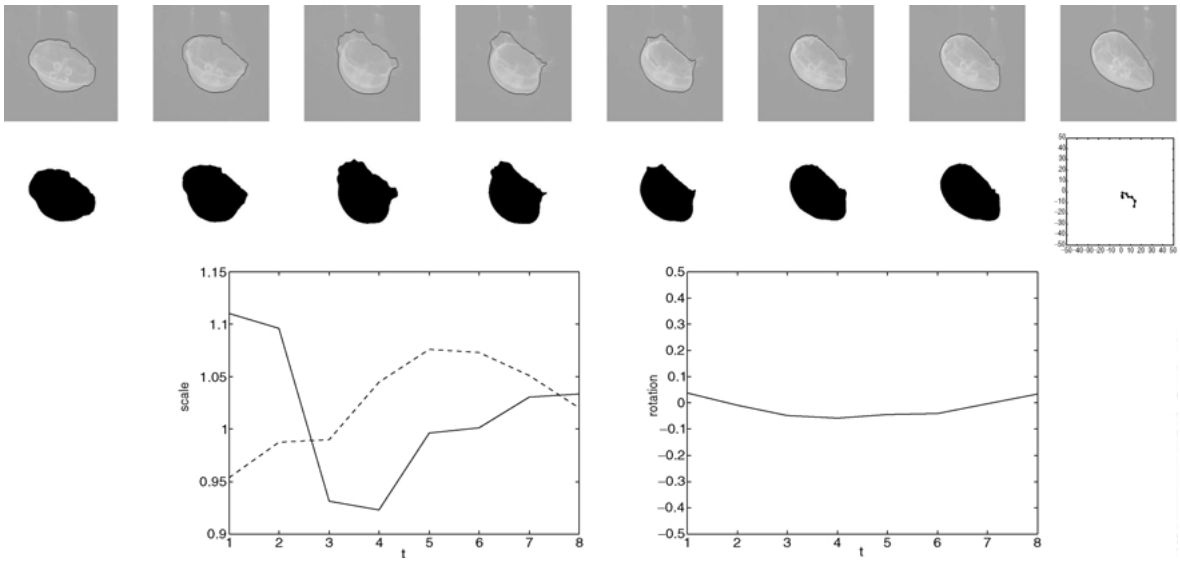

Figure 6. Jellyfish. Affine registration (top), moving average and affine motion (bottom) for the jellyfish in Fig. 1. Last row: affine scales along x and y, and rotation about z during the sequence.

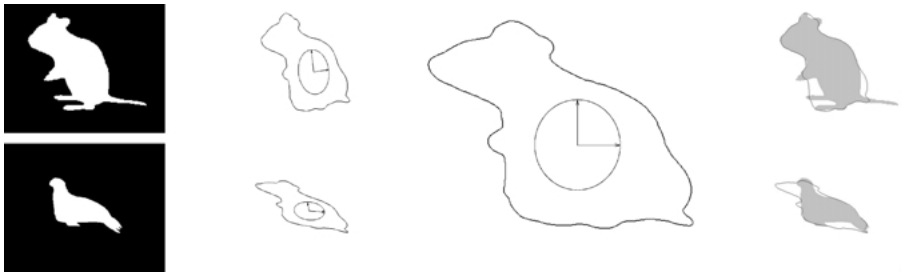

Figure 7. Registering non-equivalent shapes. Left to right: two binary images representing two different shapes; affine registration; corresponding affine shape; approximation of the original shapes using the registration of the shape average based on the set-symmetric difference. Results for the signed distance score are shown in Fig. 8.

the change of (affine) pose between the two shapes (Fig. 7). Nevertheless, it is interesting to observe how the shape average allows registering even apparently disparate shapes. Figure 8 shows a representative example from an extensive set of experiments. In some cases, the shape average contains disconnected components, in some other it includes small parts that are shared by the original dataset, whereas in others it removes parts that are not consistent among the initial shapes (e.g. the tails). Notice that our framework is not meant to capture such a wide range of variations. In particular, it does not possess a notion of "parts" and it is neither hierarchical nor compositional. In the context of non-equivalent shapes (shapes for which there is no group action mapping one exactly onto the other), the *average shape serves purely as a support to define and compute motion in a collection of images of a given deforming shape.*

Figure 9 shows the results of simutaneously segmenting and computing the average motion and registration for 4 images from a database of magnetic resonance images of the corpus callosum.

Finally, Fig. 10 shows an application of the same technique to simultaneously register and average two 3D surfaces. In particular, two 3D models in different poses are shown. Our algorithm can be used to register the surfaces and average them, thus providing a natural framework to integrate surface and volume data.

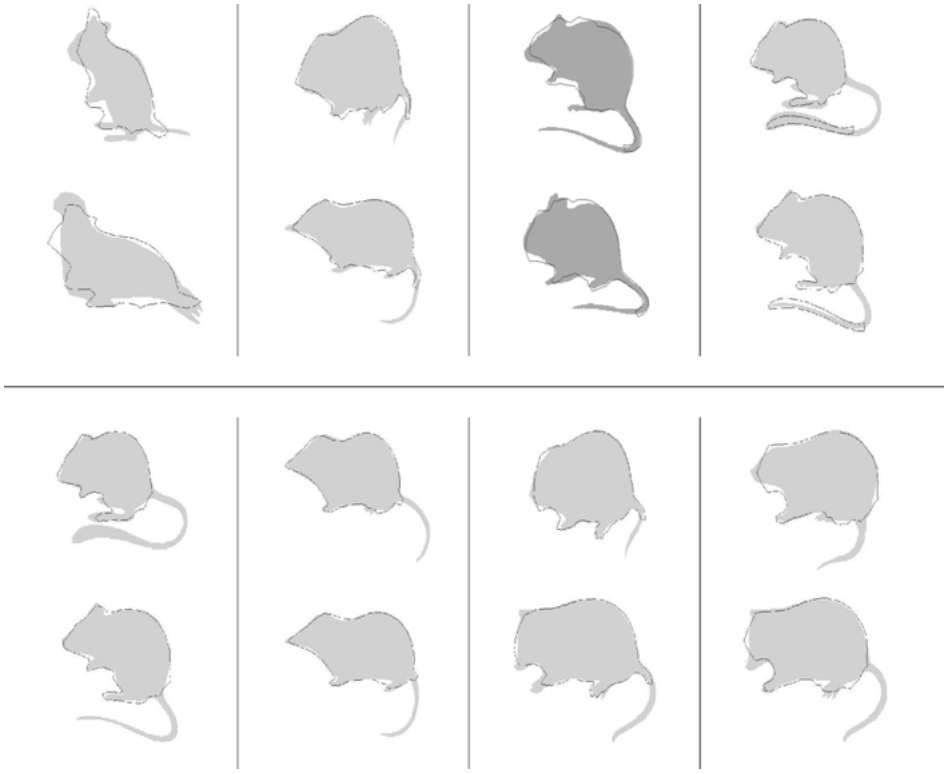

Figure 8. Biological shapes. For the signed distance score, we show the original shape with the affine shape average registered and superimposed. It is interesting to notice that in some cases the average shape is disconnected.

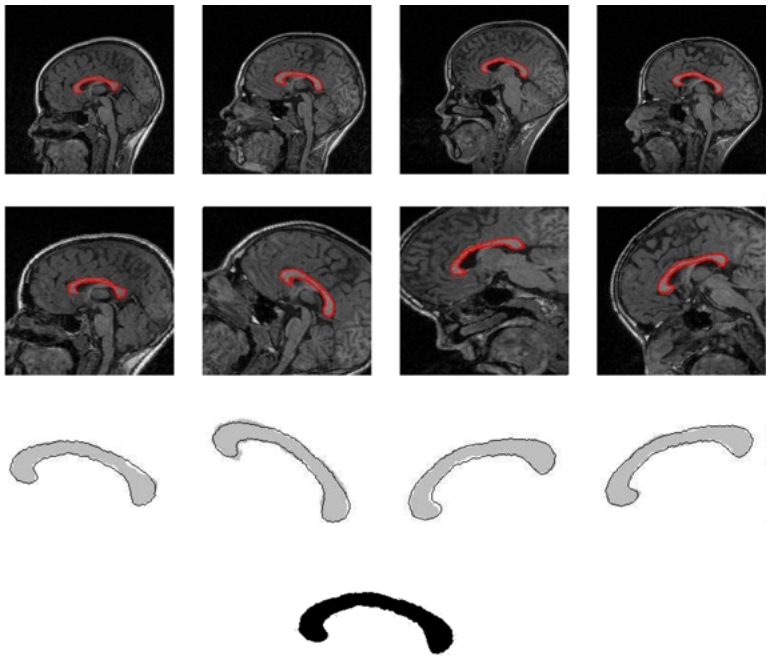

Figure 9. Corpus Callosum. (top row) a collection of (MR) images from different patients (courtesy of N. Dutta and A. Jain (2001), further translated, rotated and distorted to emphasize their misalignment, alignment and (bottom) average template corresponding to the affine group.

Figure 10. 3D Averaging and registration. (left) two images of 3D models in different poses (center) registered average (right) affine average. Note that the original 3D surfaces are not equivalent. The technique presented allows "stitching" and registering different 3D models in a natural way.

Appendix A: Rodrigues' Formula

We describe Rodrigues' formula for the case of $G = SE(3)$. The cases of $SO(3)$, $SE(2)$, $SO(2)$ follow directly as a special case. Each element (rigid motion) g is represented as a matrix:

$$g = \begin{bmatrix} R & T \\ 0 & 1 \end{bmatrix} \Bigg| \; T \in \mathbb{R}^3, \; R \in SO(3).$$

The group operations in $SE(3)$ coincide with the group operations of $\mathbb{GL}(4)$, so that the composition of rigid motions may be represented as a matrix multiplication: $g_1 \circ g_2 = G_1 G_2$. The tangent space at the origin of $SE(3)$ has the structure of a Lie algebra, and is called $se(3)$. Elements of $se(3)$ are called "twists," and may be represented in so-called "Plücker coordinates" as

$$\widehat{\xi} \doteq \dot{g} g^{-1} = \begin{bmatrix} \widehat{\omega} & v \\ 0 & 0 \end{bmatrix}, \quad \text{where } v \in \mathbb{R}^3$$

and

$$\widehat{\omega} \doteq \begin{bmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{bmatrix}$$

belongs to the Lie algebra of the skew-symmetric matrices $so(3) \doteq \{S \mid S^T = -S\}$, which is isomorphic to $\mathbb{R}^3$ via $\widehat{\omega} \leftrightarrow [\omega_1 \; \omega_2 \; \omega_3]^T \in \mathbb{R}^3$. An explicit expression for the exponential map on $SE(3)$ is given by

$$\begin{bmatrix} R & T \\ 0 & 1 \end{bmatrix} = \exp\left(\begin{bmatrix} \widehat{\omega} & v \\ 0 & 0 \end{bmatrix} \right)$$

where

$$R \doteq e^{\widehat{\omega}} = I + \frac{\widehat{\omega}}{\|\omega\|} \sin(\|\omega\|) + \frac{\widehat{\omega}^2}{\|\omega\|^2}(1 - \cos(\|\omega\|)) \tag{25}$$

$$T \doteq \frac{1}{\|\omega\|}[(I - e^{\widehat{\omega}})\widehat{\omega} + \omega\omega^T]v. \tag{26}$$

The exponential map may be inverted locally for computing v and ω from R and T when $\|\omega\| \in (0\ ,\ \pi)$. In the case $\|\omega\| = 0$, the exponential map is defined simply by

$$R \doteq I \tag{27}$$

$$T \doteq v. \tag{28}$$

Note that the exponential map, together with the isomorphism of $so(3)$ with $\mathbb{R}^3$, gives a local coordinate parametrization of $SE(3)$, which is called the canonical exponential representation. The case of $SE(2)$ can be derived simply as a special case of $SE(3)$.

Acknowledgments

We wish to thank S. Belongie for providing us with test data and suggestions and B. Kimia for discussions and references. This research is supported by NSF IIS-9876145, ECS-0200511, CCR-0133736 and IIS-0208197, ONR N00014-02-1-0720, NIH Pre-NPEBC, and Intel 8029.

Notes

1. An anonymous reviewer suggested that the recent work of Davies et al. (2001) also addresses shapes with no natural landmarks. The work of Camion and Younes (2001) was also suggested, although it appeared after our manuscript was submitted.
2. Local in this context is intended in the space of deformation functions h, rather than local to the particular object, for instance when the deformation only affects a "part" of the object.
3. The rationale behind this score is that one wants to make the integral of the signed distance transform of one contour positive as possible outside the other contour, and as negative as possible inside.
4. The "widehat" notation $\widehat{\ }$, which indicates a lifting to the Lie algebra, should not be confused with the "hat" $\hat{\ }$, which indicates an estimated quantity.
5. The identity of the group $SE(3)$ is $e = (R, T) = (O, I)$, and should obviously not be confused with the base of the natural exponential.
6. Here we use the term distance informally, since we do not require that it satisfies the triangular inequality. The term pseudo-distance would be more appropriate.

7. For instance, consider a point $\mathbf{x}(t)$ moving on the plane according to a simple linear dynamics, observed through a "noisy" measurement $\mathbf{y}(t)$:

$$\begin{cases} \mathbf{x}(t) = A_1\mathbf{x}(t-1) + A_2\mathbf{x}(t-2) + \cdots \\ \qquad\qquad + A_k\mathbf{x}(t-k) + \mathbf{v}(t) \\ \mathbf{y}(t) = C\mathbf{x}(t) + \mathbf{w}(t). \end{cases} \tag{22}$$

This model describes a dynamical system where the state $\mathbf{x}$ can be interpreted as the (autoregressive) moving average of $\mathbf{y}$. Without loss of generality one may assume that $k = 1$ since the difference equation above can always be reduced to first-order by augmenting the dimension of the state (see Ljung (1987) for details).

References

Alvarez, L., Guichard, F., Lions, P.L., and Morel, J.M. 1993. Axioms and fundamental equations of image processing. *Arch. Rational Mechanics*, 123.

Alvarez, L. and Morel, J.M. 1994. Morphological approach to multiscale analysis: From principles to equations. In *Geometric-Driven Diffusion in Computer Vision*, B.M. ter Haar Romeny (Ed.).

Alvarez, L., Weickert, J., and Sanchez, J. 1999. A scale-space approach to nonlocal optical flow calculations. In *ScaleSpace'99*, pp. 235–246.

Arnold, V.I. 1978. *Mathematical Methods of Classical Mechanics*. Springer Verlag.

Azencott, R., Coldefy, F., and Younes, L. 1996. A distance for elastic matching in object recognition. *Proc. 13th Intl. Conf. on Patt. Recog*, 1:687–691.

Belongie, S., Malik, J., and Puzicha, J. 2001. Matching shapes. In *Proc. of the IEEE Intl. Conf. on Computer Vision*.

Bereziat, D., Herlin, I., and Younes, L. 1997. Motion detection in meteorological images sequences: Two methods and their comparison. In *Proc. of the SPIE*.

Blake, A. and Isard, M. 1998. *Active Contours*. Springer Verlag.

Carne, T.K. 1990. The geometry of shape spaces. *Proc. of the London Math. Soc.*, 3(61):407–432.

Chui, H. and Rangarajan, A. 2000. A new algorithm for non-rigid point matching. In *Proc. of the IEEE Intl. Conf. on Comp. Vis. and Patt. Recog.*, pp. 44–51.

Davies, R., Cootes, T., and Taylor, C. 2001. A minimum description length approach to statistical shape modelling. In *17th Conference on Information Processing in Medical Imaging*, pp. 50–63.

Dutta, N. and Jain, A. 2001. Corpus callosum shape analysis: A comparative study of group differences associated with dyslexia, gender and handedness. In *MMBIA* submitted.

Fischler, M. and Elschlager, R. 1973. The representation and matching of pictorial structures. *IEEE Transactions on Computers*, 22(1):67–92.

Giblin, P. 1977. *Graphs, Surfaces and Homology*. Chapman and Hall.

Grenander, U. 1993. *General Pattern Theory*. Oxford University Press.

Grenander, U. and Miller, M.I. 1994. Representation of knowledge in complex systems. *J. Roy. Statist. Soc. Ser. B*, 56:549–603.

Jackway, P.T. and Deriche, R. 1996. Scale-space properties of the multiscale morphological dilationerosion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 18(1):38–51.

Kendall, D.G. 1984. Shape manifolds, procrustean metrics and complex projective spaces. *Bull. London Math. Soc.*, 16.

Kimia, B., Tannebaum, A., and Zucker, S. 1995. Shapes, shocks, and deformations i: The components of two-dimensional shape and the reaction-diffusion space. *Int'l J. Computer Vision*, 15:189–224.

Kimmel, R. 1997. Intrinsic scale space for images on surfaces: The geodesic curvature flow. In *Lecture Notes in Computer Science: First International Conference on Scale-Space Theory in Computer Vision*.

Kimmel, R. and Bruckstein, A. 1995. Tracking level sets by level sets: A method for solving the shape from shading problem. *Computer Vision, Graphics and Image Understanding*, (62)1:47–58.

Kimmel, R., Kiryati, N., and Bruckstein, A.M. 1998. Multivalued distance maps for motion planning on surfaces with moving obstacles. *IEEE Trans. Robot. & Autom.*, 14(3):427–435.

Koenderink, J.J. 1990. *Solid Shape*. MIT Press.

Lades, M., Borbruggen, C., Buhmann, J., Lange, J., von der Malsburg, C., Wurtz, R., and Konen, W. 1993. Distortion invariant object rcognition in the dynamic link architecture. *IEEE Trans. on Computers*, 42(3):300–311.

Le, H. and Kendall, D.G. 1993. The riemannian structure of euclidean shape spaces: A novel environment for statistics. *The Annals of Statistics*, 21(3):1225–1271.

Leventon, M., Grimson, E., and Faugeras, O. 2000. Statistical shape influence in geodesic active contours. In *Proc. IEEE Conference Comp. Vision and Patt. Recognition*.

Ljung, L. 1987. *System Identification: Theory for the User*. Prentice Hall.

Malladi, R., Kimmel, R., Adalsteinsson, D., Caselles, V., Sapiro, G., and Sethian, J.A. 1996. A geometric approach to segmentation and analysis of 3d medical images. In *Proc. Mathematical Methods in Biomedical Image Analysis Workshop*, pp. 21–22.

Malladi, R., Sethian, J.A., and Vemuri, B.C. 1995. Shape modeling with front propagation: A level set approach. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 17(2):158–175.

Mardia, K.V. and Dryden, I.L. 1989. Shape distributions for landmark data. *Adv. Appl. Prob.*, 21(4):742–755.

Matheron, G. 1975. *Random Sets and Integral Geometry*. Wiley.

Miller, M.I. and Younes, L. 1999. Group action, diffeomorphism and matching: A general framework. In *Proc. of SCTV*.

Mumford, D. 1991. Mathematical theories of shape: Do they model perception? In *Geometric Methods in Computer Vision*, vol. 1570, pp. 2–10.

Osher, S. and Sethian, J. 1988. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi equations. *J. of Comp. Physics*, 79:12–49.

Paragios, N. and Deriche, R. 2000. Geodesic active contours and level sets for the detection and tracking of moving objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(3):266–280.

Samson, C., Blanc-Feraud, L., Aubert, G., and Zerubia, J. 1999. A level set model for image classification. In *International Conference on Scale-Space Theories in Computer Vision*, pp. 306–317.

Sebastian, T.B., Crisco, J.J., Klein, P.N., and Kimia, B.B. 2000. Constructing 2D curve atlases. In *Proceedings of Mathematical Methods in Biomedical Image Analysis*, pp. 70–77.

ter Haar Romeny, B., Florack, L., Koenderink, J., and Viergever, M. (Eds.). 1997. Scale-space theory in computer vision. In *Lecture Notes in Computer Science*, vol. 1252, Springer Verlag.

Thom, R. 1975. *Structural Stability and Morphogenesis*. Benjamin: Reading.

Thompson, D.W. 1917. *On Growth and Form*. Dover.

Thompson, P. and Toga, A.W. 1996. A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans. Med. Imaging*, 15(4):402–417.

Veltkamp, R.C. and Hagedoorn, M. 1999. State of the art in shape matching. Technical Report UU-CS-1999-27, University of Utrecht.

Yezzi, A. and Soatto, S. 2001. Stereoscopic segmentation. In *Proc. of the Intl. Conf. on Computer Vision*, pp. 59–66.

Yezzi, A., Zollei, L. and Kapur, T. 2001. A variational approach to joint segmentation and registration. In *Proc. IEEE Conf. on Comp. Vision and Pattern Recogn*.

Younes, L. 1998. Computable elastic distances between shapes. *SIAM J. of Appl. Math*., 58(2):565–586.

Yuille, A. 1991. Deformable templates for face recognition. *J. of Cognitive Neurosci.*, 3(1):59–70.

Zhu, S., Lee, T., and Yuille, A. 1995. Region competition: Unifying snakes, region growing, energy/bayes/mdl for multi-band image segmentation. In *Int. Conf. on Computer Vision*, pp. 416–423.