DEFORMOTION
Deforming Motion, Shape Average and the Joint Registration and Segmentation of Images*

Stefano Soatto Anthony J. Yezzi

University of California, Los Angeles – CA 90095, email soatto@ucla.edu
Georgia Institute of Technology, Atlanta – GA 30332, email ayezzi@ece.gatech.edu

Keywords: Shape, segmentation, motion, registration.

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 involve the numerical integration of partial differential equations, which we address within the framework of level set methods.

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

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 [4]. 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

*A short version of this manuscript has been submitted to the ECCV on Nov. 23, 2001.

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 issue.

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 capturing the consequences of scribbling notes outdoor on a windy day. 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 making 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 Sect. 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 Sect. 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" (Sect. 4). We then extend these results from geometric shapes to images, resulting in their simultaneous approximation and registration in Sect. 5. Finally, in Sect. 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" [38], the work of Matheron on "Stochastic Sets" [30] as well as that of Thom, Giblin and others [37, 13].

In statistics, the study of "Shape Spaces" was championed by Kendall, Mardia and Carne among others [17, 24, 9, 29]. 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 under the aegis 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, 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 [40] for a recent survey. A common approach consists of matching collections points organized in graphs or trees (e.g. [23, 12]). Belongie et al. [6] 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 [10, 11]. Koenderink [22] 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 [32].

"Deformable Templates," pioneered by Grenander [14], 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 [44]. There, the notion of "motion" (or "alignment" or "registration") coincides with that of deformation, and there is no clear distinction between the two [7]. Grenander's work sparked a current that has been particularly successful in the analysis of medical images, for instance [15]. 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 [5, 43, 31, 20, 35], but quite sophisticated theories that encompass perceptually relevant aspects, can be formalized in terms of the properties of the evolution of PDEs (e.g. [21]). The work of Kimia et al. [18] 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. [18] 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 [28, 39, 27]. Although most of the ideas are develop in a deterministic setting, many can be transposed to a probabilistic context (e.g. Zhu et al. [45]). 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 in a common scale-space and then matching 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 scale-space representation at some level 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 [16, 36, 19, 2, 3, 1] and references therein. The "alignment," or "registration," of curves has also been used to define a notion of "shape average" by several authors (see [25] 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. [25] 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, and the "registration" cannot be adapted as the

original contours evolve towards their average. Similarly, [41] 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 [42], 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 [34] 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 [8] 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*[1]. 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 (Sect. 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 (Sect. 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 (Sect. 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 [33], 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.

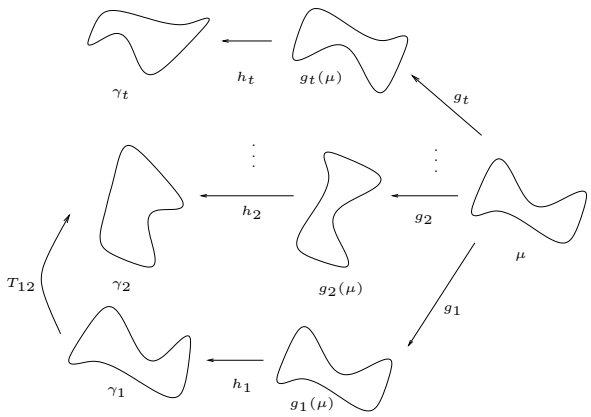

Figure 2: *A model (commutative diagram) of a deforming contour.*

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

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

It comprises the action of a group $g \in G$ (e.g. $G = SE(2)$) and a more general transformation h that belongs to a pre-defined class $\mathcal{H}$ (for instance diffeomorphisms). The deformation h is not arbitrary, but depends upon another "shape" μ, defined in such a way that

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

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 $g = g_i g_j^{-1}$ and h is a transformation that depends on h_i, h_j and μ. 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 $\sum_{t=1}^{n} E(h_t)$ subject to $\gamma_t = h_t \circ g_t(\mu)$. Note that all that matter in the cost of h_t are the "shapes" before and after the transformation, $\mu_i \doteq g_i(\mu)$ and γ_i, so that 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.

Definition 1 *Let $\gamma_1, \dots, \gamma_n$ be smooth boundaries of closed subsets of a differentiable manifold embedded in $\mathbb{R}^N$, 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 $\hat{g}_1, \dots, \hat{g}_n$ is a* **motion** *undergone by γ_i, $i = 1 \dots n$ if there exists a pre-shape $\hat{\mu}$ such that*

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

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.*

[1]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.

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 $T(g \circ g_0, h \circ g_0) = T(g, h) \ \forall \ g_0$). 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 $\gamma_1, \dots, \gamma_n$, 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, 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 Def. 1 we have purposefully avoided to use the article* "the" *for the minimizing value of the group action $\hat{g}_t$. 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 a particular element of the group G, or to an entire subgroup. 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 Sect. 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 and principal modes of variation.

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 cost functional for the deformations h_1, h_2 either the set-symmetric difference Δ of their interior (the union minus the intersection), or what we call the signed distance transform score[2] ψ

$$\psi(\mu, \gamma) \doteq \int_{\bar{\mu}} \zeta(\gamma) d\mathbf{x} + \int_{\bar{\gamma}} \zeta(\gamma) d\mathbf{x} \tag{5}$$

where $\bar{\mu}$ denotes the interior of the contour μ and ζ is the signed distance function of the contour γ; $d\mathbf{x}$ 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

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

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 Def. 1, we can write the problem above as an unconstrained minimization

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

[2]The rationale behind this score is that one wants to make the signed distance function as positive as possible outside the contour to be matched, and as negative as possible inside. This score can be interpreted as a weighted Monge-Kantorovic functional where the mass of a curve is weighted by its distance from the boundary.

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 functional above is just a particular case 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 binary images, we have $f_{in} = (\chi_\gamma - 1)^2$ and $f_{out} = \chi_\gamma^2$. 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

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

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[3]

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

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 App. 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

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

[3]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.

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 absent 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 $\nabla_G \phi = \widehat{\left(\frac{\partial \phi}{\partial \xi}\right)} g$. 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 $\int_{g(\mu)} \left\langle f(\mathbf{x}) \frac{\partial g}{\partial \xi_i} \circ g^{-1}(\mathbf{x}), N \right\rangle ds = \int_{\mu} \left\langle f \circ g(\mathbf{x}) \frac{\partial g}{\partial \xi_i}, g_* N \right\rangle ds$ 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 contour and motion are then updated in an alternating minimization fashion where motion is updated according to

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

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 $\frac{\partial \phi}{\partial \xi_i} = \int_{g(\gamma_{1,2})} \left\langle \frac{\partial g(\mathbf{x})}{\partial \xi_i}, f_{\{in,out\}}(g(\mathbf{x}), \gamma_{1,2}) J(g_* T) \right\rangle ds$. The contour $\hat{\mu}$ evolves according to

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

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 (6) suggests defining the distance[4] 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

$$d_{\Delta}(\gamma_1, \gamma_2) \doteq \int \chi_{\hat{\mu}} \chi_{\gamma_1} + \chi_{\hat{g}(\hat{\mu})} \chi_{\gamma_2} d\mathbf{x} \tag{20}$$

[4]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.

and for the signed distance transform we have

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

In either case, given two contours, a gradient flow algorithm based on Eq. (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.

3.5 Principal modes of variation and "Harmonic Embedding"

Once a collection of contours γ_i has been aligned via $\hat{g}_i^{-1}(\gamma_i)$, and their average $\hat{\mu}$ has been characterized, one may want to define principal modes of variation of the collection. Leventon et al. [25] have proposed performing principal component analysis of the contours represented as their signed distance function. Unfortunately, however, signed distance functions do not form a linear vector space, and therefore the interpretation of the resulting principal directions is somewhat unclear. In this section we propose a different repreentation that is *linear* by construction, so that linear analysis (e.g. principal components) and interpolation can be performed, is invariant to the group G, is *intrinsic*, i.e. it does not depend on a particular parametrization, uses a *minimal* number of basis elements given a training set and is characterized by a certain *regularity*.

The intrinsic requirement suggests representing shape as the levelset of the steady-state solution of a PDE. The linearity requirements suggests restricting the attention to linear PDEs. The regularity requirement suggests avoiding first-order PDEs. Therefore, the requirements described suggest representing shape as the zero level set of the steady-state solution of a second-order PDE. A natural choice of Laplace's equation. In order to proceed, let Γ_{-1} be the boundary of the region enclosed by each (registered) contour and Γ_{+1} the boundary of the region that encloses all contours. The region between Γ_{-1} and Γ_{+1} defines an annulus $\mathcal{A}$. We are interested in characterizing the collection of contours in $\mathcal{A}$. Laplace's squation leads to a choice of harmonic functions that satisfy the following constraints:

$$\begin{aligned} u(x) &= 0 \ \forall \ x \in \gamma & (22) \\ \Delta u(x) &= 0 \ \forall \ x \in \mathcal{A}. & (23) \end{aligned}$$

Of all functions that satisfy these constraints we want a unique choice to be the representative of the shape of γ. We choose u in such a way that it will have a prescribed value on one of the boundaries of the annulus, for instance

$$u(x) = -1 \ \forall \ x \in \Gamma_{-1} \tag{24}$$

which will result in a value function at the other boundary:

$$u(x) = f(x) > 0 \ \forall \ x \in \Gamma_{+1} \tag{25}$$

we call u the *harmonic embedding* of γ.

The representation of the shape of a given contour γ using its harmonic embedding u form a linear vector space. Therefore, defining and computing distances between shapes, averages of shapes and projections onto subspaces spanned by a collection of shapes can be performed in a straightforward fashion in terms of their harmonic embedding. For instance, projections, principal components, orthogonal complements can be defined and computed in a straightforward fashion.

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," 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,"* similarly to what is done in time series analysis[5]. 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

$$\begin{cases} \mu(t+1) = g(t)\mu(t) \\ \gamma(t) = h(\mu(t)) \end{cases} \tag{27}$$

where $\mu(t)$ represents the moving average of order $k = 1$. The procedure described in Sect. 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 (27) is represented as the model parameter. Note that the procedure in Sect. 3 simultaneously estimates the state $\mu(t)$ and identifies the parameters $g(t)$ of the model (27). 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.

$$\begin{cases} g(t+1) = e^{\widehat{\xi}(t)} g(t) \\ \mu(t+1) = g(t)\mu(t) \\ \gamma(t) = h(\mu(t)) \end{cases} \tag{28}$$

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 Eq. (17)-(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 algorith 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: $g = g_2 g_1^{-1}$.

If instead of considering the images in Fig. 7 as binary images that represent the contours, we consider them as grayscale images, then the procedure outlined, for the case where the score is computed using the set-symmetric difference, provides a way to simultaneously jointly segment the two images and register them. This idea is illustrated in Fig. 9 for true grayscale (magnetic resonance) images of brain sections. Therefore, this procedure can be used to simultaneously segment a collection of uncalibrated images, thus extending the approach of [41].

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

$$\begin{cases} \mathbf{x}(t) = A_1\mathbf{x}(t-1) + A_2\mathbf{x}(t-2) + \cdots + A_k\mathbf{x}(t-k) + \mathbf{v}(t) \\ \mathbf{y}(t) = C\mathbf{x}(t) + \mathbf{w}(t). \end{cases} \tag{26}$$

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 [26] for details).

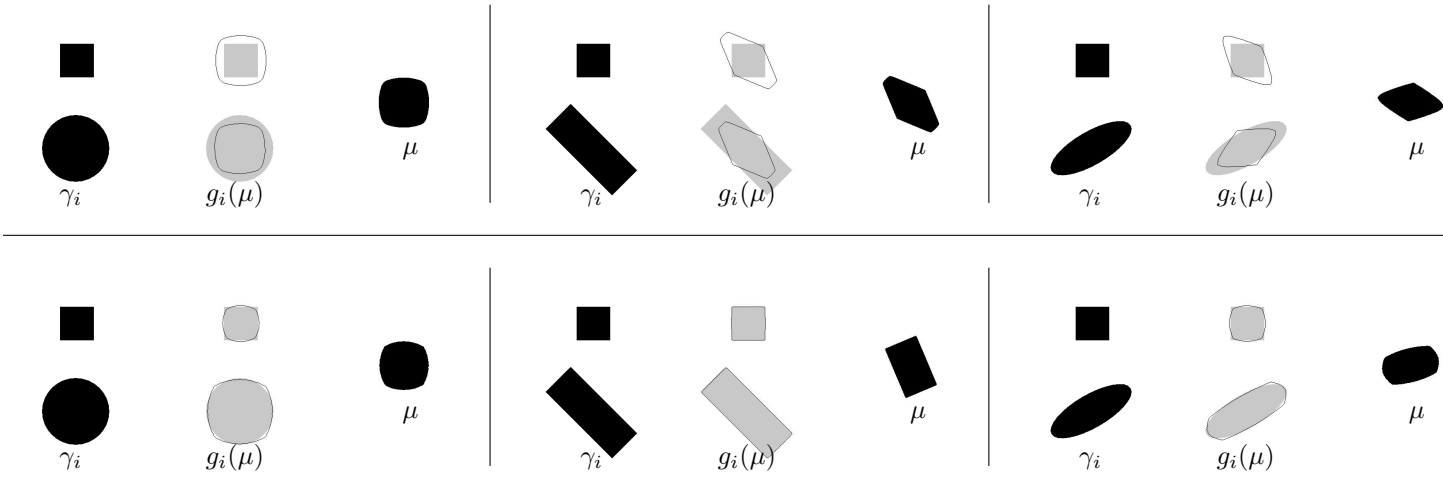

Figure 3: **Euclidean (top) vs. affine (bottom) registration and average.** *For each pair of objects γ_1, γ_2, the registration $g_1(\mu), g_2(\mu)$ relative to the Euclidean motion and affine motion is shown, together with the Euclidean average and affine average μ. Note that the affine average can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

6 Experiments

Fig. 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 to simultaneously capture the square and the rectangle, whereas the Euclidean average cannot be registered to either one, and is therefore only an approximation.

Fig. 4 compares the effect of choosing the signed distance transform 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.

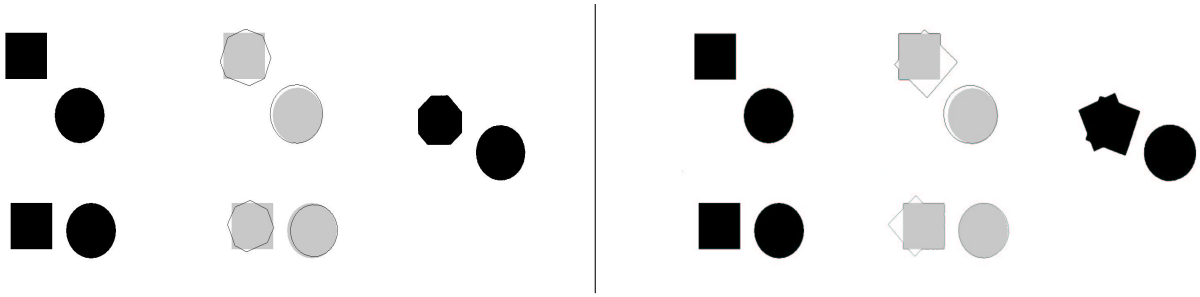

Figure 4: **Signed distance transform score (left) vs. set-symmetric difference (right).** *Original contours (γ_1 on the top, γ_2 on the bottom), registered shape $g_i(\mu)$ and shape average μ. Note that the original objects are not connected, but are composed by a circle and a square. The choice of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

Fig. 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.

Fig. 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 the change of (affine) pose between the two shapes

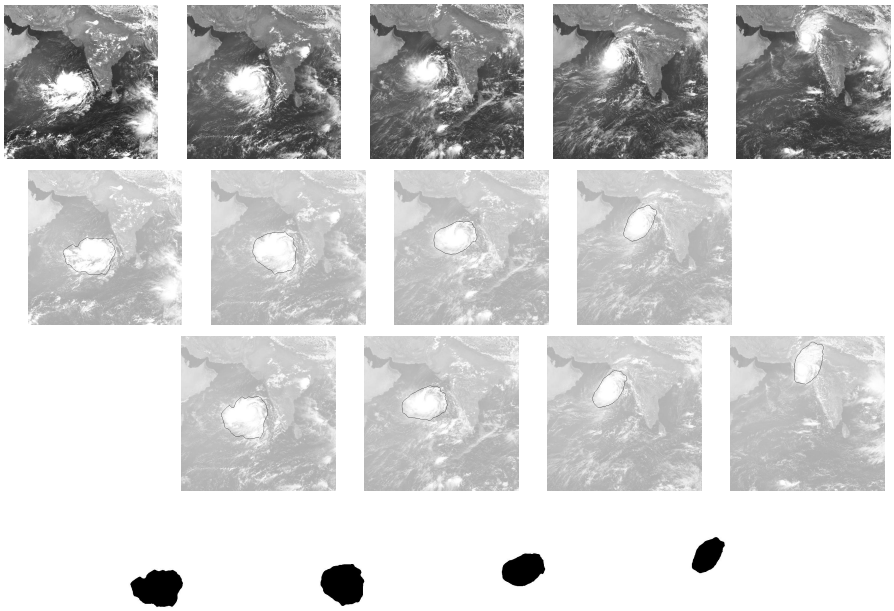

Figure 5: **Storm** *(first row) a collection of images from EUMETSAT ©2001, affine motion of the storm based on two adjacent time instances, (bottom) moving average of order 1.*

(Fig. 7). Nevertheless, it is interesting to observe how the shape average allows registering even apparently disparate shapes. Fig. 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.*

Fig. 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.

Acknowledgements

We wish to thank S. Belongie for providing us with test data and suggestions. This research is supported in part by NSF grant IIS-9876145, ARO grant DAAD19-99-1-0139 and Intel grant 8029.

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:

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

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

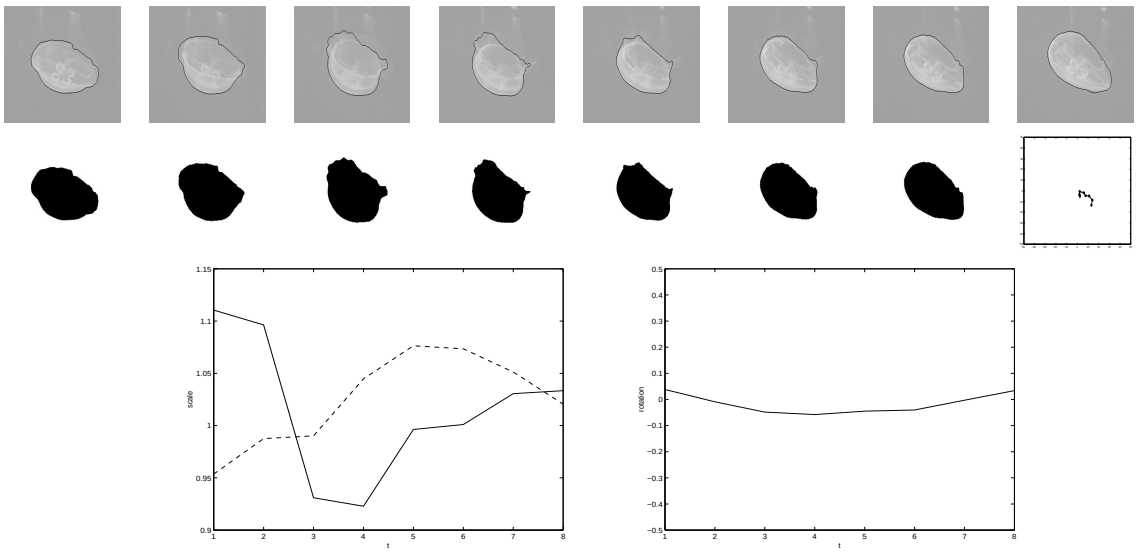

Figure 6: **Jellyfish.** *Affine registration (top), moving average and affine motion (bottom) for the jellyfish in Fig. 1. Last row: affine scales along x and y, and rotation about z during the sequence.*

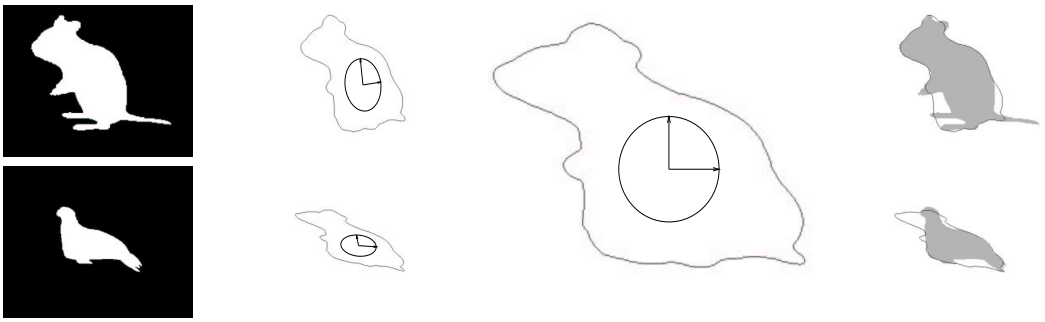

Figure 7: **Registering non-equivalent shapes.** *Left to right: two binary images representing two different shapes; affine registration; corresponding affine shape; approximation of the original shapes using the registration of the shape average based on the set-symmetric difference. Results for the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

$\widehat{\omega} \doteq \begin{bmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{bmatrix}$ belongs to the Lie algebra of the skew-symmetric matrices $so(3) \doteq \{S | S^T = -S\}$, 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

$$R \doteq e^{\widehat{\omega}} = I + \frac{\widehat{\omega}}{\|\omega\|}\sin(\|\omega\|) + \frac{\widehat{\omega}^2}{\|\omega\|^2}\Big(1 - \cos(\|\omega\|)\Big) \tag{29}$$

$$T \doteq \frac{1}{\|\omega\|}\left[\left(I - e^{\widehat{\omega}}\right)\widehat{\omega} + \omega\omega^T\right]v. \tag{30}$$

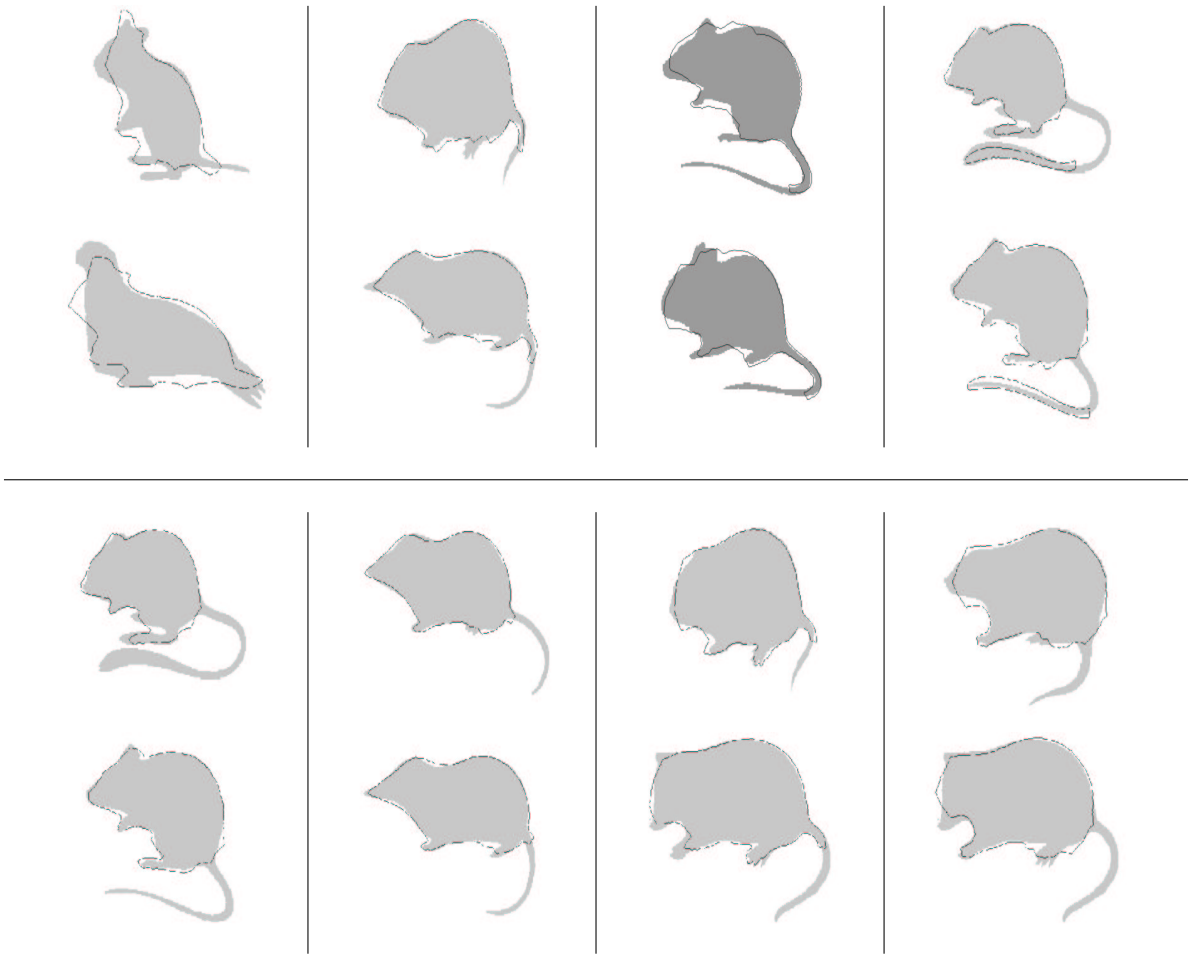

Figure 8: **Biological shapes** *For the signed distance transform score, we show the original shape with the affine shape average registered and superimposed. It is interesting to notice that different "parts" are captured in the average only if they are consistent in the two shapes being matched and, in some cases, the average shape is disconnected.*

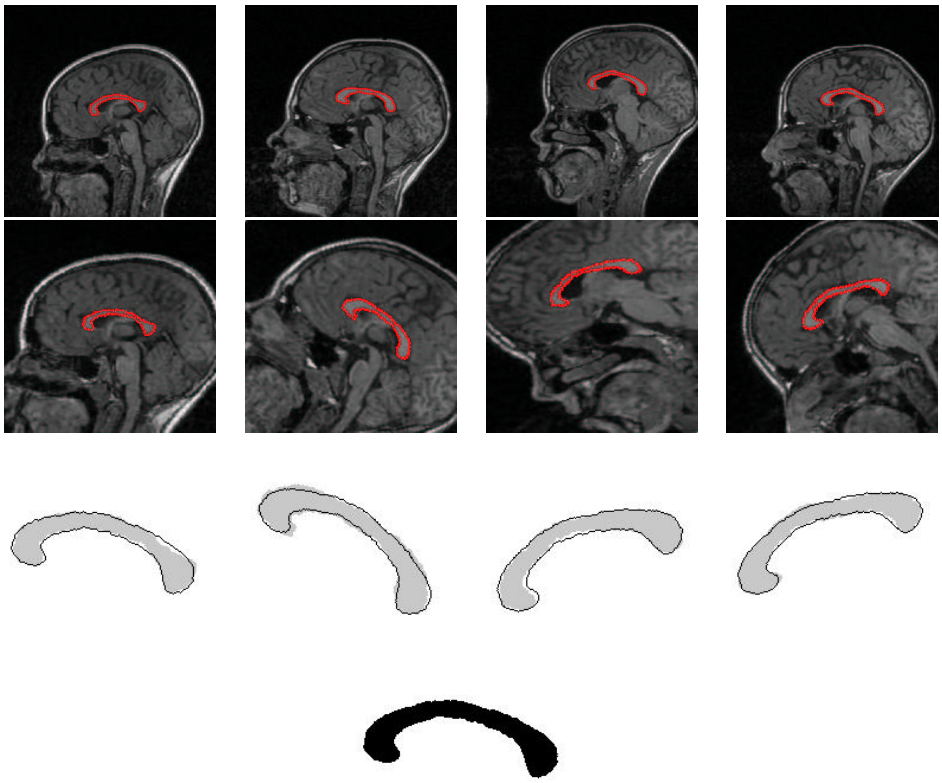

Figure 9: **Corpus Callosum** *(top row) a collection of (MR) images from different patients (courtesy of N. Dutta and A. Jain [11]), further translated, rotated and distorted to emphasize their misalignment, alignment and (bottom) average template corresponding to the affine group.*

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

$$R \doteq I \tag{31}$$

$$T \doteq v. \tag{32}$$

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 $SE(3)$.

References

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

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

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

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

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

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

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

[8] A. Blake and M. Isard. *Active contours*. Springer Verlag, 1998.

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

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

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

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

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

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

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

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

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

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

[19] R. Kimmel. 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*, 1997.

[20] R. Kimmel and A. Bruckstein. 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, 1995.

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

[22] J. J. Koenderink. *Solid Shape*. MIT Press, 1990.

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

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

[25] M. Leventon, E. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours, 2000.

[26] L. Ljung. *System Identification: theory for the user*. Prentice Hall, 1987.

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

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

[29] K. V. Mardia and I. L. Dryden. Shape distributions for landmark data. *Adv. appl. prob.*, 21(4):742–755, 1989.

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

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

[32] D. Mumford. Mathematical theories of shape: do they model perception? In *In Geometric methods in computer vision, volume 1570*, pages 2–10, 1991.

[33] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi equations. *J. of Comp. Physics*, 79:12–49, 1988.

[34] N. Paragios and R. Deriche. 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, 2000.

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

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

[37] R. Thom. *Structural Stability and Morphogenesis*. Benjamin; Reading, 1975.

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

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

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

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

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

[43] L. Younes. Computable elastic distances between shapes. *SIAM J. of Appl. Math.*, 1998.

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

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